

July 19, 2026

To the ENG2003 Technical Review Panel:

Enclosed is the final technical report, “AI-Augmented Oncology: Leveraging Deep Learning for Earlier Cancer Detection,” prepared as the Phase 4 (Revise, Finalize and Reflect) deliverable for Term Project 1 in ENG2003 — Effective Engineering Communication. This report was commissioned as part of the coursework, following an agreement to submit a final proposal-style technical report by July 19, 2026.

The report addresses physician-scientist Eric Topol’s grand challenge of “digitizing human beings” — the difficulty of converting the growing volume of individual health data into timely, actionable insight. It narrows that challenge to a single application: using deep learning applied to medical imaging to detect cancer earlier than conventional radiology and pathology allow. The report explains how convolutional neural networks function in this context. It reviews recent trial evidence from mammography and colonoscopy screening, and it discusses the strengths, limitations, and adoption barriers of this technology from a biomedical and computer engineering perspective. It closes with a recommendation on the technology’s near-term role in oncology practice.

This version incorporates revisions made in response to Phase 3 peer feedback. A Revision Summary describing that feedback and the resulting changes is included as an appendix at the end of the report. Questions or requests for clarification regarding the content or structure of this report are welcome. Thank you for your consideration.

Respectfully submitted,
Parsa Abbasian

AI-Augmented Oncology:

Leveraging Deep Learning for Earlier Cancer Detection

A Technical Report on Applying Convolutional Neural Networks to Medical Imaging for Early-Stage Cancer Diagnosis

Prepared for: ENG2003 — Effective Engineering Communication

Term Project 1, Phase 4 (Finalize)

Discipline: Biomedical / Computer Engineering — Medical Imaging & Health Informatics

Date: July 19, 2026

Author: Parsa Abbasian

Executive Summary

Physician-scientist Eric Topol has identified the challenge of “digitizing human beings” — aggregating and interpreting the growing volume of individual health data — as one of the defining grand challenges of the 21st century. This report examines one concrete application of that challenge: applying deep learning, specifically convolutional neural networks (CNNs), to medical imaging to detect cancer earlier and more consistently than conventional radiology and pathology practice.

Cancer remains a leading cause of death worldwide, with outcomes depending heavily on both cancer type and stage at diagnosis. Recent trial evidence indicates that AI-assisted image analysis can meaningfully narrow this early-detection gap. In the 2026 MASAI trial, the first large randomized controlled trial of AI in breast cancer screening, AI-supported mammography reading increased cancer detection by 29% and reduced screen-reading workload by 44%, without increasing false positives. In colonoscopy, AI-assisted polyp-detection systems have raised adenoma detection rates from roughly 33–47% to 41–59% across multiple randomized and quasi-randomized trials. These gains are occurring against a backdrop of rapid regulatory growth: the U.S. FDA had cleared over 1,300 AI/ML-enabled medical devices by December 2025, nearly 80% of them in radiology.

At the same time, this technology has real limitations. Performance still lags conventional readers in some conditions such as dense breast tissue, models can encode dataset bias, and broader adoption raises questions of data privacy, cost, and workforce impact that are not yet fully resolved.

This report concludes that AI-augmented imaging should be adopted as a clinician-support tool — improving consistency and reducing workload — rather than a replacement for human judgment. It recommends continued investment paired with bias auditing and transparency safeguards as the technology matures.

Table of Contents

Executive Summary.....	3
Table of Contents.....	4
List of Figures	5
List of Tables	5
1.0 Introduction	6
2.0 Background	6
3.0 How Deep Learning Enables Earlier Cancer Detection	7
4.0 Current Performance Across Imaging Modalities	7
4.1 Mammography Screening.....	8
4.2 Colonoscopy.....	8
4.3 Broader Applications Across Modalities	9
5.0 Strengths, Limitations, and Adoption Barriers.....	10
6.0 Discussion.....	10
7.0 Conclusion	11
References	12
AI Use Disclosure Statement.....	13
Appendix A: Revision Summary	15

List of Figures

Figure 1: The AI-Augmented Cancer Detection Pipeline — Section 4.0

Figure 2: AI-Assisted vs. Conventional Adenoma Detection Rates — Section 4.2

List of Tables

Table 1: Selected Performance Outcomes of AI-Assisted vs. Conventional Detection, by Modality — Section 4.0

1.0 Introduction

Physician-scientist Eric Topol, of the Scripps Translational Science Institute, has named the capacity to “digitize human beings” as one of the defining grand challenges of the 21st century [9]. By this he means using biosensors, DNA sequencing, and imaging to build a continuous, individual medical profile. Topol argues that the real difficulty is not collecting this data but processing it. He proposes aggregating the resulting terabytes of health information with algorithms, and particularly deep learning, to build what he calls a “virtual medical coach.” Such a system would give people instantaneous feedback on their health and help prevent illness before it progresses [9].

This report narrows Topol’s broad challenge to a single, high-stakes application: cancer. In this disease, the gap between an available diagnosis and the disease’s biological head start often separates a curable illness from a terminal one. The proposed technological response is artificial intelligence, specifically deep learning and convolutional neural networks (CNNs), applied to medical imaging. Relevant modalities include mammography, CT, MRI, colonoscopy, and digital pathology. The goal is to detect cancerous and pre-cancerous findings earlier and more consistently than conventional radiology and pathology practice alone. The discipline of inquiry is biomedical and computer engineering, specifically medical imaging and health informatics.

The purpose of this report is to evaluate whether, and under what conditions, deep-learning-based image analysis can meaningfully close this early-detection gap. Its contents include background on the scale of the cancer burden and why detection timing matters, an explanation of how CNN-based detection systems operate, evidence of real-world performance drawn from three imaging modalities, a discussion of the technology’s strengths and adoption barriers, and a concluding recommendation. This report does not evaluate AI applications downstream of detection, such as drug discovery or treatment planning, nor does it assess the cost economics of clinical implementation.

2.0 Background

Cancer is among the leading causes of death worldwide. The World Health Organization estimates that roughly one in five people will develop cancer during their lifetime, with over 20 million new cases diagnosed annually and nearly 10 million deaths [5]. Critically, outcomes vary enormously by cancer type and, above all, by the stage at which the disease is caught: U.S. Surveillance, Epidemiology, and End Results (SEER) data compiled by the National Cancer Institute show five-year survival ranging from 99% for thyroid cancer down to just 13% for pancreatic cancer, a gap driven largely by how early each cancer type is typically detected [3]. Early detection is therefore not a peripheral concern in oncology; it is frequently the single largest determinant of whether a patient survives.

Conventional cancer screening relies on trained radiologists and pathologists visually reviewing images: mammograms, CT and MRI scans, colonoscopy video, and pathology slides. This process faces well-documented limits. Radiologists reading high volumes of mammograms under time pressure can miss subtle, early-stage lesions, particularly in patients with dense breast tissue where tumors and normal tissue appear visually similar. Colonoscopists can miss small or flat polyps that are easy to overlook during

a live procedure. Pathologists reviewing digital slides face a similarly high cognitive load across large tissue samples. None of these limitations reflect a lack of skill; they reflect the inherent difficulty of sustained, high-stakes visual pattern recognition performed at scale by a finite number of trained specialists.

This is precisely the kind of problem Topol’s grand challenge anticipates: an abundance of available data (in this case, medical images) that outpaces the human capacity to process it consistently and exhaustively. Topol specifically identifies deep learning as the key enabling technology for closing this gap, arguing that algorithms trained on large datasets can supplement, and in some tasks match or exceed, human visual pattern recognition [9]. The following sections examine whether this promise holds up under recent clinical evidence, beginning with an explanation of how these systems work.

3.0 How Deep Learning Enables Earlier Cancer Detection

The technology underlying most AI-based cancer detection systems is the convolutional neural network (CNN), a type of deep learning model designed specifically to process image data. A CNN learns to recognize patterns — edges, textures, shapes, and eventually complex structures such as tumor margins or tissue abnormalities — by passing an image through successive layers of small filters that each detect increasingly abstract features. Rather than being explicitly programmed with rules for what a cancerous lesion looks like, the network is trained on large sets of medical images that have already been labeled by expert radiologists or pathologists as malignant, benign, or normal. Over many training examples, the model adjusts its internal parameters until it can reliably associate certain image patterns with a diagnosis [1].

Once trained, a CNN-based system processes a new image and outputs a probability score, and often a visual heatmap, indicating the likelihood and location of malignancy. Critically, in nearly all current clinical deployments, this output is not used to replace the radiologist or pathologist but to support them: the AI system functions as either a “second reader,” flagging cases for a human expert to review, or as a triage tool, prioritizing the most concerning cases for faster review. This human-in-the-loop design is a deliberate response to the limitations of the technology, discussed further in Section 5.0.

The regulatory infrastructure surrounding these tools has expanded rapidly in step with their clinical use. As of late 2025, the U.S. Food and Drug Administration’s list of authorized AI/ML-enabled medical devices had grown to more than 1,300 entries, of which 1,039 (nearly 80%) were radiology-specific tools; in 2024 alone, the FDA authorized 168 new Class II machine-learning-enabled devices, the large majority cleared through the streamlined 510(k) pathway on the basis of similarity to existing predicate devices [8]. This growth signals that CNN-based image analysis has moved well past the research stage and is being treated by regulators as a maturing, clinically deployable technology, particularly within radiology and, by extension, oncologic imaging. The next section reviews what this technology has actually achieved in recent clinical trials.

4.0 Current Performance Across Imaging Modalities

Figure 1 illustrates the general processing pipeline described in Section 3.0, from image acquisition through to a clinician’s final diagnostic decision. This structure underlies each of the imaging applications discussed below.



Figure 1: The AI-Augmented Cancer Detection Pipeline

4.1 Mammography Screening

The strongest evidence to date comes from the Mammography Screening with Artificial Intelligence (MASAI) trial, a randomized controlled study of more than 105,000 women and the first large-scale randomized trial of AI in breast cancer screening. Its 2026 results, published in *The Lancet*, found that AI-supported single reading identified 29% more cancers than the conventional double-reading process, while achieving a non-inferior 12% reduction in the rate of interval cancers (cancers diagnosed between scheduled screenings) and a 44% reduction in radiologists’ screen-reading workload, all without an increase in false positives [6]. However, a comprehensive 2025 systematic review notes that AI’s advantage is not uniform: radiologists retain higher sensitivity than standalone AI systems in patients with dense breast tissue, where AI systems tend to show higher specificity but can under-call subtle findings [10]. This nuance supports the “second reader” deployment model described in Section 3.0 rather than full automation.

4.2 Colonoscopy

In colorectal cancer screening, AI-assisted polyp-detection (“computer-aided detection,” or CAdE) systems have shown consistent gains across multiple trials. A 2025 multicenter quasi-randomized controlled trial found that AI assistance raised the adenoma detection rate from 46.6% to 59.1% [7]. A broader 2023 systematic review and meta-analysis of randomized clinical trials pooled these results across studies and found a similarly significant effect, with adenoma detection rates of 41.4% in AI-assisted colonoscopies compared to 33% in standard colonoscopies, corresponding to roughly double the odds of detecting an adenoma or polyp during the procedure [11]. These gains were most pronounced for smaller, more easily missed polyps, precisely the category most likely to be caught before progressing to cancer. Figure 2 compares the adenoma detection rates from these two colonoscopy studies.




Figure 2: AI-Assisted vs. Conventional Adenoma Detection Rates (Sources [7], [11])

4.3 Broader Applications Across Modalities

Beyond mammography and colonoscopy, a January 2025 review analyzed AI applications across ten cancer types spanning medical imaging, genomic analysis, and digital pathology, and concluded that AI-assisted approaches consistently improved diagnostic accuracy, sensitivity, and speed relative to traditional radiology and pathology workflows, while also enabling newer non-invasive detection methods such as liquid biopsy [1]. These broader findings, however, rest largely on narrative review rather than the large randomized trials that anchor the mammography and colonoscopy evidence above; the following discussion therefore weighs these headline performance gains against the operational realities — uneven evidence quality, dataset bias, and adoption barriers — that determine whether they translate into clinical practice. Table 1 summarizes the quantitative outcomes discussed above.

Modality	AI-Assisted Outcome	Conventional Outcome	Source
Mammography (MASAI RCT, 2026)	29% higher cancer detection rate; 44% lower reading workload; no increase in false positives	Standard double reading, no AI	[6]
Mammography, dense breast tissue	Higher specificity; lower sensitivity than human readers in dense tissue	Higher sensitivity in dense tissue	[10]
Colonoscopy (quasi-RCT, 2025)	Adenoma detection rate: 59.1%	Adenoma detection rate: 46.6%	[7]
Colonoscopy (RCT meta-analysis, 2023)	Pooled ADR: 41.4% (approx. 2x odds of detection)	Pooled ADR: 33%	[11]
FDA-cleared AI/ML devices, all specialties	1,039 radiology-specific tools cleared by Dec. 2025 (approx. 80% of all AI device clearances); 168 new devices authorized in 2024 alone	N/A — regulatory volume, not a diagnostic comparison	[8]

Table 1: Selected Performance Outcomes of AI-Assisted vs. Conventional Detection, by Modality

Together, these results indicate that CNN-based detection is not a speculative future technology but one already producing measurable, reproducible improvements in two of the highest-volume cancer

screening programs in medicine, with evidence of similar promise extending across a much broader range of cancer types.

5.0 Strengths, Limitations, and Adoption Barriers

The clearest strength of AI-assisted detection is consistency at scale. Unlike a human reader, a trained model applies the same evaluation criteria to the one-thousandth image as to the first, without fatigue. This directly addresses the workload and burnout pressures on radiologists and pathologists documented in Section 2.0. As the MASAI and colonoscopy results demonstrate, it can improve both detection rates and reading throughput simultaneously. This scalability also has implications for democratizing access to cancer screening. It is particularly valuable in regions with a shortage of specialist radiologists and pathologists, and it directly advances the equity-oriented vision underlying Topol’s original challenge [9].

At the same time, the evidence reviewed in Section 4.0 also points to real limitations. Performance is not uniform across patient populations. AI systems currently show reduced sensitivity relative to human readers in dense breast tissue, a condition affecting a substantial share of screening populations [10]. More broadly, CNN-based models are trained on specific datasets and can encode the biases present in that data. They perform less reliably when applied to patient populations, imaging equipment, or disease presentations that differ from their training data. Deep learning models are also frequently described as “black boxes.” They produce a probability score without an easily interpretable explanation, which complicates clinical trust and accountability when a model’s judgment conflicts with a human reader’s.

Beyond these technical limitations, a critical 2024 review of AI in healthcare highlights adoption barriers that are organizational and ethical rather than purely technical. These include data privacy risks associated with aggregating sensitive medical imaging data, algorithmic bias leading to unequal quality of care across patient groups, the high cost of implementing and maintaining these systems, and legitimate concern about the impact on radiology and pathology workforces [2]. That review argues that these governance and equity questions must be addressed alongside technical performance. Otherwise, the benefits of this technology risk being distributed unevenly, undermining the egalitarian premise of “digitizing human beings” for the benefit of all patients rather than only those served by well-resourced health systems [2], [9].

6.0 Discussion

The evidence assembled in this report suggests that AI-augmented medical imaging is beginning to deliver on Topol’s grand challenge, but only under specific conditions. Where evidence is strongest — mammography and colonoscopy — the technology has moved from small retrospective studies to large, multicenter randomized trials, the gold standard of clinical evidence, and the results are consistent: meaningful gains in detection rates paired with reduced clinician workload. This is a materially different evidentiary position than existed even a few years ago and represents genuine progress on closing the early-detection gap described in Section 2.0.

However, the rapid growth in FDA clearances documented in Section 3.0 is concentrated overwhelmingly in radiology, and the underlying trial evidence is similarly concentrated in a handful of high-volume

screening applications. Cancer types and imaging modalities outside this narrow band, along with digital pathology more broadly, have comparatively thinner clinical trial evidence despite the promising narrative-review findings discussed in Section 4.3 [1], [4]. This unevenness matters: a technology proven in mammography screening does not automatically generalize to, for example, rare pediatric cancers or resource-limited pathology labs, and treating it as though it does risks the exact equity failure that the critical-perspective source in Section 5.0 warns against [2].

The recurring pattern across the strongest studies is that AI performs best as a decision-support tool that augments, rather than replaces, a human expert — improving consistency and catching cases a fatigued or time-pressured reader might miss, while a clinician retains final diagnostic authority and continues to outperform AI in specific edge cases such as dense breast tissue. This “second reader” model appears to be the most defensible and, based on current regulatory trends, the most likely near-term deployment pathway for this technology in oncologic imaging.

7.0 Conclusion

This report examined whether deep learning applied to medical imaging can meaningfully advance Eric Topol’s grand challenge of digitizing human beings for earlier, more consistent cancer detection. The evidence indicates that it can, within a defined scope: randomized trial data from the 2026 MASAI mammography study and multiple colonoscopy trials show statistically significant improvements in cancer and adenoma detection rates alongside reduced clinician workload, and regulatory clearances of AI-enabled imaging devices have grown rapidly since 2024. At the same time, the technology’s performance is not yet uniform across patient populations or imaging modalities, and its broader adoption raises unresolved questions of bias, data privacy, and cost.

This report recommends continued clinical adoption of AI-assisted imaging as a decision-support tool that augments, rather than replaces, radiologists and pathologists, paired with ongoing bias auditing across diverse patient populations and continued transparency reporting of the kind the FDA has begun to require of newly cleared devices. Further research and trial evidence are particularly needed for cancer types and modalities, such as digital pathology, where the current evidence base remains narrower than in mammography and colonoscopy. With these safeguards, AI-augmented imaging represents a credible and increasingly well-evidenced path toward the earlier, more equitable cancer detection that Topol’s original challenge envisions.

References

- [1] M. Aftab, F. Mehmood, C. Zhang, A. Nadeem, Z. Dong, Y. Jiang, and K. Liu, "AI in oncology: Transforming cancer detection through machine learning and deep learning applications," arXiv:2501.15489, Jan. 2025.
- [2] M. A. Rahman, E. Victoros, J. Ernest, R. Davis, Y. Shanjana, and M. R. Islam, "Impact of artificial intelligence (AI) technology in healthcare sector: A critical evaluation of both sides of the coin," *Clinical Pathology*, vol. 17, 2024, doi: 10.1177/2632010X241226887.
- [3] National Cancer Institute, "Cancer Statistics," NCI, 2024. [Online]. Available: <https://www.cancer.gov/about-cancer/understanding/statistics>
- [4] C. Fang, P. Zhou, X. Zhang, Y. He, and Q. Yang, "Artificial intelligence in oncology drug development and management: A precision medicine perspective," *Frontiers in Oncology*, 2025, doi: 10.3389/fonc.2025.1609827.
- [5] World Health Organization, "Cancer," WHO Fact Sheets, 2024. [Online]. Available: <https://www.who.int/news-room/fact-sheets/detail/cancer>
- [6] K. Lång et al., "Interval cancer, sensitivity, and specificity comparing AI-supported mammography screening with standard double reading without AI in the MASAI study: A randomised, controlled, non-inferiority, single-blinded, population-based, screening-accuracy trial," *The Lancet*, vol. 407, pp. 505–514, Jan. 2026.
- [7] R. M. B. Lagström, K. B. Bräuner, J. Bielik, A. W. Rosen, J. G. Crone, I. Gögenur, and M. Bulut, "Improvement in adenoma detection rate by artificial intelligence-assisted colonoscopy: Multicenter quasi-randomized controlled trial," *Endoscopy International Open*, vol. 13, a25215169, Feb. 2025, doi: 10.1055/a-2521-5169.
- [8] B. Almarie, L. F. Gonzalez-Gonzalez, L. A. dos Santos Barbosa, A. Lutz, U. Grosse, et al., "Machine learning-enabled medical devices authorized by the US Food and Drug Administration in 2024: Regulatory characteristics, predicate lineage, and transparency reporting," *Biomedicines*, vol. 13, no. 12, art. 3005, Dec. 2025.
- [9] B. Lufkin, "The 50 grand challenges facing our world," BBC Future, 2017. [Online]. Available: <http://www.bbc.com/future>
- [10] E. Abu Abeeelh and Z. Abuabeileh, "Screening mammography and artificial intelligence: A comprehensive systematic review," *Cureus*, vol. 17, no. 2, e79353, Feb. 2025, doi: 10.7759/cureus.79353.
- [11] S. Lou, F. Du, W. Song, Y. Xia, X. Yue, D. Yang, et al., "Artificial intelligence for colorectal neoplasia detection during colonoscopy: A systematic review and meta-analysis of randomized clinical trials," *eClinicalMedicine*, vol. 66, art. 102341, 2023.

AI Use Disclosure Statement

In accordance with the ENG2003 course policy, I am disclosing the use of generative AI in the preparation of this report. The core research and foundational writing are entirely my own work. I independently researched, selected, and read all primary sources to draft the Phase 1 annotated bibliography, which included crafting the summaries, evaluations, and reflections. Furthermore, I fully wrote the initial Phase 2 draft report. My research specifically focused on Eric Topol's challenge of "digitizing human beings," and I independently developed the angle exploring the integration of Convolutional Neural Networks (CNNs) in medical imaging technologies.

I utilized Gemini during Phase 4 of the assignment to assist with the revision and formatting processes. Specifically, I used the AI to synthesize and organize the PeerScholar feedback provided by three peers into a more actionable checklist. I also prompted the tool to identify excessively long sentences in Sections 1.0 and 5.0 and suggest ways to break them up for better readability. Additionally, the AI assisted in generating the two visual figures included in the report, adjusting the data placement within Table 1 to ensure proper alignment, and applying standard formatting to the Table of Contents, reference list, and overall document layout.

It is important to emphasize that I rigorously verified all AI-generated suggestions before incorporating them. I manually cross-checked every statistic and technical claim against my original peer-reviewed sources, ensuring exact accuracy for metrics such as the MASAI 29% figure, the colonoscopy Adenoma Detection Rate (ADR) numbers, and the FDA device approval counts. I did not blindly accept the AI's synthesis of my peer feedback; I independently reviewed the raw comments and pushed back on suggestions that did not align with my intent. I also reviewed the generated figures for accuracy and heavily edited the AI's sentence-restructuring suggestions to ensure the phrasing remained in my authentic voice, outright rejecting several changes that altered the technical nuance of my arguments.

I confirm that no data, citations, or references were fabricated. The foundational arguments, critical analysis, and final conclusions presented in this report represent my own genuine understanding of the selected literature.

I received peer feedback that several sentences in my report rely too heavily on semicolons and em-dashes, making them hard to follow. Here is the text from Section 1.0 and Section 5.0. Can you identify the longest, most complex sentence? 

I can certainly help you refine these sections. Your peers make a valid point: technical writing is often most effective when complex ideas are delivered in clear, digestible pieces.

I have reviewed the text from Sections 1.0 and 5.0 and identified a few dense sentences that carry a heavy cognitive load. Here is an analysis of those sentences and suggestions on how to break them into shorter, self-contained statements while preserving your technical nuance.

[DOCX](#)

Section 1.0: Introduction

Original Long Sentence:

"Its contents include background on the scale of the cancer burden and why detection timing matters, an explanation of how CNN-based detection systems operate, evidence of real-world performance drawn from three imaging modalities, a discussion of the technology's strengths and adoption barriers, and a concluding recommendation." [DOCX](#)

Analysis:

This sentence contains a five-part list separated by commas, which forces the reader to hold a lot of structural information in their working memory.

Suggested Revision (Splitting into two sentences):

"This report first establishes the scale of the cancer burden and the critical importance of detection timing. It then explains the operation of CNN-based systems, reviews their real-world performance across three imaging modalities, and evaluates current adoption barriers before offering a concluding recommendation."

Appendix A: Revision Summary

This appendix summarizes the peer feedback received on the Phase 2 draft through PeerScholar and the revisions made in response. Three peers reviewed the draft, with rubric scores of 14/16, 15/16, and 14/16.

1. Sentence length and readability (Sections 1.0 and 5.0). Peer 1 noted that several sentences were too long and relied heavily on semicolons and em-dashes, making dense passages hard to follow, and specifically pointed to Sections 1.0 and 5.0. I agreed with this feedback. In the revised report I broke the longest compound sentences in both sections into shorter, self-contained statements and removed several semicolons and em-dashes in favour of full stops. The technical content is unchanged; only the sentence structure was simplified to improve readability.

2. Section 4.3 feeling too general (transition and framing). Peer 1 observed that the report’s evidence is concentrated on mammography and colonoscopy, which made Section 4.3 (“Broader Applications”) and the Discussion feel more general than the rest of the report, and suggested adding a transition at the end of Section 4.3 to bridge into the operational realities. I agreed. I added a transition sentence at the end of Section 4.3 that explicitly notes the broader findings rest on narrative review rather than randomized trials, and that signals the shift into the discussion of evidence quality, bias, and adoption barriers. This makes the change in scope deliberate rather than abrupt.

3. Table 1 data placement. Peer 1 pointed out that in the final row of Table 1, the FDA device-count figures were placed under the “Conventional Outcome” column, which was out of place because that row reports regulatory volume rather than a diagnostic comparison. I agreed. I moved the “168 new devices authorized in 2024” figure into the AI-Assisted Outcome column alongside the related clearance data, and replaced the Conventional Outcome cell with a note clarifying that this row reports regulatory volume, not a head-to-head diagnostic result.

4. Table of Contents / page-numbering mismatch. Peer 2 caught that the Table of Contents listed an “AI Use Disclosure Statement” on page 12, but the draft ended on page 11, so the referenced section did not exist. Rather than simply deleting the entry, I resolved the mismatch by writing and adding the AI Use Disclosure Statement as an actual section, which the course’s AI policy expects in any case. The Table of Contents is now generated automatically so that all entries and page numbers stay consistent with the final document.

5. Reconciling and weighing the feedback. The three reviewers did not give contradictory advice; their comments were complementary. Peer 2 and Peer 3 both rated the report highly overall, with Peer 2 praising the balanced treatment of the technology’s strengths and limitations and Peer 3 describing it as a good technical report overall. Because Peer 3 gave no specific, actionable points, I could not act on that review beyond the general encouragement to keep the existing structure. I therefore prioritized Peer 1’s and Peer 2’s concrete, actionable suggestions, all of which I accepted. Beyond the peer feedback, I also completed a full proofreading and formatting pass across the whole report, checking citation consistency, figure and table captions, and page numbering, in line with the guidance to improve the report holistically rather than only where peers commented.

6. Sections left unedited, added visual, and overall reflection. Sections 2.0 (Background), 3.0 (How Deep Learning Enables Earlier Cancer Detection), 6.0 (Discussion), and 7.0 (Conclusion) were left substantively unchanged. No peer flagged these sections, and on my own re-reading they already met the technical-report criteria: the background establishes the cancer burden and screening limits with cited evidence, the technology explanation is clear and audience-appropriate, and the discussion and conclusion follow logically from the evidence without overstating it. Editing them further risked weakening passages the reviewers had identified as strengths. In addition to the peer-driven edits, I added a new figure (Figure 2) charting the colonoscopy adenoma-detection results already reported in Table 1, so that the strongest paired evidence is presented visually as well as in text. Comparing the final report to the original draft, the core argument, evidence base, and structure are unchanged, but the final version reads more clearly at the sentence level, presents its data more accurately in Table 1 and the new figure, and is internally consistent between the table of contents and the actual document.