International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Artificial Intelligence and Robotics as an Integrated "Detect–Target–Confirm" Loop for Precision Oncology: Mapping and Manipulating Cancer Cells in Localized Anatomical Regions

[1]**Dr. Chaitanya Kumar M V**, Professor in ECE SOET, Christ University, Mysore Road, Bangalore 560074 chaitanya.kumar@christuniversity.in,
[2]**Dr.Ashoka S. B**, Associate Professor, Department of Computer Science, Maharani Cluster University, Bangalore India. dr.ashoksbcs@gmail.com
[3]**Dr.Niranjan R Chougala**, Professor and Head Computer Science & Technology, RRIT (VTU Autonomous) Bangalore. drnrchougala@gmail.com, niranjan.chougala@rrit.ac.in
[4]**Mr. Shashank H S**, Assistant Professor, Department of Computer Science, Maharani Cluster University, Bangalore India. shashankshettyhs@gmail.com

ABSTRACT

Cancer care increasingly depends on the ability to find malignant cells with sub-millimetre precision and to act on them without damaging the healthy tissue around them. This paper examines how artificial intelligence (AI) and robotic systems are converging into a closed-loop clinical workflow that can be usefully described as "detect, target, and confirm." In this model, AI operates as the cognitive layer, parsing radiological images, histopathology slides, and molecular biomarker data to localize tumours and characterize their genetic and phenotypic traits, while robotic platforms — ranging from macro-scale surgical manipulators to magnetically steered micro- and nanorobots — translate that intelligence into physical action inside a defined region of the body. We review the architecture of this loop across three stages: detection through deep-learning image analysis and multi-omic data fusion; targeting through robot-assisted surgery, image-guided biopsy, and actively steered drug carriers; and confirmation through intraoperative sensing, real-time imaging feedback, and biomarker verification that closes the control loop and

lets the system adapt mid-procedure. Evidence from recent robotic-surgical studies, nanorobotic simulations, and deep-learning segmentation challenges is synthesised to illustrate the current capability and maturity of each stage. We also discuss unresolved challenges, including the generalizability of imaging models across patient populations and scanners, the difficulty of navigating soft, deformable tissue at the microscale, regulatory and safety questions, and the burden of integrating imaging, robotics, and biological data into a single decision pipeline. We conclude that although no single commercial system yet performs the entire loop autonomously, its individual components are maturing quickly enough that a genuinely closed detect-target-confirm system for localized cancer care is a realistic near-term objective rather than a speculative one.

Keywords: artificial intelligence, medical robotics; oncology, tumour detection; targeted therapy, nanorobotics, robotic surgery, precision medicine

1. Introduction

Cancer remains one of the leading causes of death worldwide, and much of the difficulty in treating it comes down to a problem of precision. A tumour is rarely a uniform mass; it is a patchwork of cell populations with different genetic drivers, growth rates, metabolic characteristics, and sensitivities to treatment, embedded within healthy tissue that must be preserved [2,3]. This biological and spatial heterogeneity makes cancer treatment fundamentally different from many other clinical interventions. The objective is not simply to identify the presence of a tumour, but to determine its precise boundaries, understand its biological characteristics, select an appropriate therapeutic strategy, deliver that intervention accurately, and verify whether the intended treatment effect has actually been achieved. Failure at any one of these stages can influence treatment outcomes, increase complications, or result in residual disease.

Clinicians have historically had to rely heavily on their own visual interpretation of medical images, their own hands to guide a needle or a scalpel, and their own judgement to decide, often during a procedure, whether a lesion has been adequately sampled or removed. Medical imaging modalities such as computed tomography (CT),

magnetic resonance imaging (MRI), positron emission tomography (PET), ultrasound, and digital pathology provide enormous amounts of information, but interpreting this information remains a demanding task. Small lesions, poorly defined tumour boundaries, anatomical deformation, imaging artefacts, and variations between patients can make accurate localisation challenging. In addition, the biological characteristics of a tumour cannot always be inferred from anatomical imaging alone. Molecular, histopathological, and multi-omic information may provide additional evidence, but integrating these heterogeneous data sources into a single clinical decision remains difficult [3,12,14,18]. Consequently, cancer care continues to face a central challenge: translating large volumes of complex diagnostic information into precise and timely physical action.

The emergence of artificial intelligence (AI), particularly machine learning and deep learning, has created new possibilities for addressing this challenge. AI systems can process large and heterogeneous datasets and identify patterns that may be difficult to recognise consistently through conventional manual analysis. In medical imaging, deep learning models have demonstrated strong capabilities in tumour detection, classification, segmentation, image registration, and treatment-response assessment [3,12,14,18]. Similar developments have occurred in digital pathology, where computational methods can analyse tissue morphology at cellular and subcellular levels and assist in identifying malignant regions and clinically relevant features. AI can also potentially integrate information from multiple modalities, including radiological images, pathological findings, genomic profiles, and clinical records, thereby providing a more comprehensive representation of tumour characteristics.

However, accurate prediction or detection alone does not solve the clinical problem. An AI model may identify a suspicious region or estimate a tumour boundary, but that information must ultimately be translated into a physical intervention. This is where medical robotics provides an important complementary capability. Modern robotic systems can provide enhanced instrument stability, precise positioning, tremor reduction, motion scaling, and access to anatomically constrained regions. Robot-assisted surgical platforms have therefore become increasingly important in minimally invasive procedures and are being investigated across different areas of oncologic surgery [1,6]. Beyond conventional surgical robots, advances in micro- and nanorobotics have introduced the possibility of manipulating therapeutic agents or miniature devices at much smaller spatial scales. Magnetically guided micro/nanorobots, for example, are being investigated for targeted drug delivery and other localized therapeutic applications [9]. These developments suggest that robotics can provide the physical precision required to convert computationally generated treatment decisions into targeted actions.

The potential value of combining AI and robotics lies not merely in placing two technologies side by side, but in establishing a continuous interaction between perception, decision-making, intervention, and verification. AI can provide the capability to interpret complex information and identify clinically meaningful targets, while robotic systems can translate those targets into controlled physical movements or localized therapeutic actions. Once an intervention has been performed, new imaging, sensing, or tissue information can be analysed by AI to determine whether the desired objective has been achieved. The result is a potentially continuous cycle in which information is acquired, interpreted, acted upon, and evaluated again.

This paper therefore frames AI-enabled medical robotics for precision oncology as a closed-loop “detect–target–confirm” framework. In the detect stage, AI-based methods analyse medical images, histopathological data, molecular information, and other clinical inputs to identify and characterise suspected tumour regions. The target stage translates these AI-derived findings into actionable spatial information, allowing robotic surgical instruments, needles, therapeutic delivery systems, or micro/nanorobotic platforms to reach the intended location with high precision. The confirm stage evaluates the outcome of the intervention using imaging, sensing, pathology, or other feedback mechanisms. The information generated during confirmation can then be used to update the system's understanding of the treatment site and, where appropriate, support further intervention. In this interpretation, AI functions as the system's cognitive or “brain” layer, whereas robotics provides the physical or “hands” layer.

The closed-loop perspective is particularly important because conventional clinical workflows are often organised as a sequence of separate stages. Diagnosis may be performed using one technology, treatment planning using another, intervention using a separate platform, and outcome assessment at a later time. Information can therefore be lost or transformed as it moves between stages. A closed-loop architecture instead seeks to

maintain a continuous relationship between what the system observes, what it decides, what it does, and what it subsequently measures. Such an approach has the potential to reduce uncertainty and improve adaptability, particularly in situations where anatomy changes during surgery or where the response to treatment cannot be predicted completely in advance.

Another important motivation for this integration is the increasing availability of multimodal medical data. Modern oncology generates information from radiology, pathology, genomics, proteomics, electronic health records, intraoperative imaging, and physiological sensors. Each data source provides a different perspective on disease. AI can potentially combine these sources to construct a richer representation of tumour location and biological behaviour. Robotics can then use this information to execute highly localized interventions. For example, an AI system could identify a tumour boundary from preoperative imaging, refine the target using intraoperative imaging, and guide a robotic instrument toward the region while continuously accounting for anatomical movement. Similarly, in targeted drug delivery, AI-based analysis could potentially assist in identifying an appropriate treatment location while a micro/nanorobotic platform performs the localized delivery. These examples illustrate how multimodal perception and physical actuation can become interconnected rather than remaining isolated technological capabilities.

Despite these advances, the integration of AI and robotics in oncology is still at an emerging stage. Most existing systems remain specialized rather than fully autonomous. An AI model may perform highly accurate segmentation without directly controlling a treatment device, while a surgical robot may provide highly precise manipulation but depend entirely on a clinician for target selection and decision-making. Similarly, micro/nanorobotic drug-delivery systems are promising but face substantial challenges in navigation, localization, biocompatibility, controllability, and clinical translation [6,9]. The absence of reliable feedback mechanisms is another important limitation. Without confirmation, a system may know where it intended to act but not whether the intervention achieved the desired biological or anatomical outcome.

The transition from open-loop assistance to closed-loop precision oncology therefore requires more than improvements in individual algorithms or robotic hardware. It requires reliable communication between perception and action, accurate spatial registration, robust uncertainty estimation, real-time sensing, safe control mechanisms, and clinically meaningful feedback. AI predictions must be sufficiently reliable to support robotic action, while robotic systems must be capable of operating within defined safety boundaries and responding appropriately to uncertainty. Furthermore, any system that influences clinical decisions must account for issues such as data variability, model generalisation, interpretability, cybersecurity, human oversight, and regulatory requirements. These considerations become particularly important when AI-generated decisions are translated directly into physical actions within the human body. From a clinical perspective, the ultimate objective is not to replace the clinician but to augment clinical capabilities. Human expertise remains essential for understanding patient-specific circumstances, interpreting uncertain evidence, managing unexpected events, and making ethical and therapeutic decisions. A well-designed AI–robotic system should therefore function as an intelligent and controllable partner, providing enhanced perception and precision while preserving appropriate human oversight. This human-in-the-loop approach may represent an important transitional pathway toward increasingly autonomous systems while maintaining safety and accountability.

The convergence of AI, advanced medical imaging, robotic surgery, and micro/nanorobotic drug delivery consequently represents an important direction for precision oncology. Rather than viewing these technologies as independent tools, they can be considered components of an integrated cyber-physical clinical system in which information continuously guides physical action and physical outcomes continuously generate new information. Such a system could eventually support more accurate tumour localisation, more selective treatment delivery, improved preservation of healthy tissue, and more immediate assessment of therapeutic effectiveness.

Against this background, this paper develops a unified detect–target–confirm framework for understanding the integration of AI with medical robotics in localized precision oncology. The framework considers AI-based perception and decision-making, robotic execution at surgical and micro/nanoscale levels, and real-time confirmation as interconnected components of a closed control loop.

2. The Detect–Target–Confirm Framework

The detect-target-confirm framework is best understood not as three separate technologies bolted together, but as a feedback loop, in the same sense that a thermostat is a loop rather than just a thermometer and a heater. AI does the sensing and the interpreting; robotics does the acting; and a confirmation step feeds fresh sensor data back into the AI layer so that the next action can be corrected if the first one was imprecise. This closed-loop structure is what distinguishes the systems discussed here from earlier generations of computer-assisted diagnosis, which stopped at producing a report for a human to act on, and from earlier generations of surgical robots, which were teleoperated extensions of a surgeon's hands with little independent perception of the tissue they were touching [1,5].

Three properties recur across implementations of this loop, regardless of whether the acting element is a full-sized surgical arm or a micron-scale carrier. First, the AI component must localize a target within a coordinate frame that the robotic component can act on — that is, detection has to be spatially registered, not just a binary yes/no classification. Second, the robotic component needs enough degrees of freedom and enough positional accuracy to reach that location without collateral damage to surrounding tissue. Third, some form of real-time or near-real-time sensing

optical, acoustic, magnetic, or biochemical — has to verify that the intervention landed where it was supposed to, so that the loop can close rather than run open.

2.1 Detect: AI as the Diagnostic Brain

The detection stage is the most mature of the three and has received the most research attention to date. Convolutional neural networks and, more recently, vision transformer architectures now underpin most automated tumour segmentation pipelines, with hybrid CNN-transformer models reporting classification accuracies above 98% and Dice similarity scores above 0.90 on public brain tumour benchmarks such as BraTS [14,18]. In head-and-neck cancer, AI-generated tumour contours produced for radiotherapy planning have been shown to match or exceed the consistency achieved between different human radiologists, a milestone because inter-observer disagreement has long been one of the practical bottlenecks in radiotherapy target definition [15].

What makes the detection stage relevant to a detect-target-confirm loop rather than a stand-alone diagnostic tool is the move toward multimodal, foundation-style models that fuse imaging with other data streams. One widely cited 2024 model trained on tens of thousands of whole-slide pathology images from multiple institutions achieved area-under-curve values near 0.97 in cross-validated cancer detection across eleven tumour types, and was explicitly designed to generalize across different scanners and patient populations rather than being tuned to a single dataset [3]. Other multimodal approaches combine radiology, clinical records, and liquid-biopsy genomic data using recurrent and convolutional architectures to improve early detection of cancers such as lung cancer, where a single imaging modality is often ambiguous on its own [3]. Reviews of AI in oncology also point to models that link tumour molecular subtype prediction directly to imaging features, effectively letting the AI layer output not just "where is the tumour" but "what kind of tumour is it," which is the information a downstream targeting system needs in order to select an appropriate intervention or drug payload [3,10,17].

This is the sense in which AI functions as the "brain" of the loop: it does not merely flag an abnormality, it produces a spatially localized, trait-annotated map of a specific anatomical region that a robotic system can subsequently act upon.

2.2 Target: Robotics as the Precision Hands

The targeting stage covers everything from full-sized robotic surgical arms operating inside a body cavity to micro- and nanoscale carriers navigating the bloodstream, and the appropriate platform depends heavily on the physical scale of the intervention required.

At the macro scale, robot-assisted surgical systems such as the da Vinci platform have been in clinical use for over two decades, but the addition of AI-driven guidance is what has begun to close the loop with the detection stage described above [2]. A 2025 systematic review of AI-integrated robotic cancer surgery, restricted to studies using real patient data, found consistent reports of improved resection precision and reduced operative complications when AI-based intraoperative decision support was layered onto robotic platforms, although the review also noted that the number of studies meeting strict real-patient, peer-reviewed criteria remained relatively small as of 2025 [1]. A related synthesis of 25 peer-reviewed studies published in 2024 and 2025 reported that AI-assisted robotic procedures were associated with roughly a quarter less operative time, around a 30% reduction in intraoperative complications, and measurably improved targeting accuracy during tumour resection compared with manual techniques [5]. In thoracic oncology specifically, combining AI, augmented reality, and robotic assistance has been linked to improved resection accuracy and shorter operating times relative to conventional approaches,

alongside fewer postoperative complications and better recovery of lung function [3].

At the micro and nano scale, the targeting problem changes character: instead of a surgeon-guided arm, the acting agent is a self-propelled or externally steered carrier small enough to travel through the vasculature. Magnetically actuated micro-robots coated with drug-loaded vesicles and coupled to magnetotactic bacteria for propulsion have been proposed as a way to combine active tumour-homing behaviour with a conventional chemotherapeutic payload, letting the carrier navigate toward a tumour site rather than relying purely on passive diffusion [9]. Reviews of AI-guided nanorobotics point out that passive drug-delivery systems typically achieve targeting efficiencies below 1%, and that intelligently steered carriers using magnetic or biochemical guidance have shown roughly a fivefold to twentyfold improvement in tumour drug concentration in preclinical models compared with passive nanocarriers [7,8]. Simulation work using reinforcement learning has demonstrated that a nanorobot can be trained to adjust its trajectory in response to real-time biomarker concentration readings, effectively letting the carrier "decide" how to home in on cancerous tissue rather than following a fixed path, which is a direct nanoscale analogue of the AI-guided targeting seen in robotic surgery [6,11].

Across both scales, the common thread is that the AI layer's spatial and molecular map from the detection stage is what tells the robotic layer, whether a seven-degree-of-freedom surgical arm or a magnetically driven particle, where to go and, increasingly, what specific molecular signature to look for once it gets there.

2.3 Confirm: Closing the Loop

The confirmation stage is the least standardized of the three but is arguably what distinguishes an integrated detect-target-confirm system from a simple pipeline in which AI hands off a static plan to a robot. Confirmation requires a sensing modality that operates during or immediately after the intervention and that can be fed back into the AI layer for verification or correction.

In robotic surgery, this has taken the form of AI-augmented-reality overlays that continuously re-register intraoperative imaging against the surgical field, allowing the system to flag whether the resection margin achieved matches the margin the detection-stage model recommended [3]. Systematic reviews of AI-integrated robotic oncologic surgery specifically categorise intraoperative decision support as one of the core functions distinguishing newer systems from earlier robotic platforms that offered no independent verification of surgical progress [1]. At the microscale, confirmation looks different: rather than a visual overlay, it takes the form of biomarker-triggered activation, where a nanorobot's payload is released only once local biochemical cues, such as the acidic pH characteristic of the tumour microenvironment, confirm that the carrier has actually reached malignant tissue rather than healthy tissue nearby [7,8]. In reinforcement-learning-based nanorobot simulations, this same logic is expressed algorithmically: the system continuously compares sensed biomarker concentration against an expected profile and adjusts its next move accordingly, which is functionally a real-time confirm-and-correct step rather than a one-shot targeting decision [6].

What both scales share is that confirmation is what allows the loop to be adaptive rather than open. Without it, an AI-generated tumour map is simply handed to a robot once and never checked again; with it, the system can detect a missed margin, an off-target release, or an incomplete resection and respond within the same procedure rather than only at the next follow-up scan.

3. Evidence from Recent Studies

The strength of the evidence base differs considerably between the macro-scale surgical implementation of this loop and the micro/nanoscale implementation, and it is important to be explicit about this difference when evaluating the maturity of the overall framework. Current studies demonstrate that the individual components of the detect–target–confirm architecture are progressing at different rates. AI-based detection and segmentation have reached relatively high levels of technical performance, AI-assisted robotic surgery has begun to demonstrate measurable benefits using real patient data, whereas intelligent micro/nanorobotic intervention remains predominantly at the preclinical and simulation stage.

On the detection side, recent studies provide substantial evidence that deep-learning systems can identify and spatially characterize tumour regions with high accuracy under controlled evaluation conditions. Brain tumour segmentation studies using MRI data, for example, have reported Dice similarity scores exceeding 0.90 for strong models, while some CNN-, transformer-, and hybrid architectures have achieved classification accuracies above 98% on publicly available benchmark datasets [12,14,18]. These results are important for the proposed framework because segmentation provides the spatial information required for subsequent robotic targeting. A binary classification such as "tumour present" is insufficient for robotic intervention; the system must determine where the tumour is located, estimate its

boundaries, and ideally distinguish tumour tissue from surrounding healthy structures. Recent developments in promptable and foundation-style segmentation further suggest that tumour mapping can become more adaptable to different anatomical sites and imaging conditions [16,18].

Evidence from head-and-neck tumour segmentation provides another important perspective. The 2024 HNTS-MRG challenge demonstrated continued interest in automated tumour contour generation for MRI-guided applications, where accurate target delineation is particularly important for treatment planning and image-guided intervention [15]. Such studies highlight the practical value of reducing variability between human observers. In clinical oncology, differences between experts in defining tumour boundaries can influence treatment planning, and AI-assisted contouring has the potential to provide a more consistent initial estimate that can subsequently be reviewed by clinicians. Within the detect–target–confirm architecture, this creates a natural transition from automated detection to spatially registered targeting.

The development of multimodal AI provides an additional layer of evidence for the detection stage. Modern cancer diagnosis increasingly involves the interpretation of radiological images alongside histopathological, molecular, and clinical information. Recent reviews describe AI systems capable of combining these heterogeneous sources to support cancer detection, classification, prognosis, and treatment selection [2,3,10,11,17]. A particularly important development is the movement from simply identifying the presence of cancer toward characterising tumour phenotype and molecular properties. Such information could become highly relevant to robotic intervention because the target is not necessarily defined only by its anatomical coordinates. A clinically useful target may also need to be associated with a particular molecular marker, tissue characteristic, or therapeutic vulnerability. Consequently, the detection stage is gradually evolving from a two-dimensional image-recognition problem into a multimodal tumour-mapping problem.

The evidence for the targeting stage is strongest in robot-assisted surgery. Robotic surgical platforms have already demonstrated their ability to provide precise instrument manipulation, stable motion, improved visualization, and enhanced access to anatomically constrained regions. The addition of AI-based guidance and intraoperative decision support represents an important step toward connecting these physical capabilities with computational tumour localization [1,5]. A 2025 systematic review of AI integration into robotic cancer surgery, restricted to studies involving real patient data, reported consistent indications of improved surgical precision and reduced complications when AI-based capabilities were incorporated into robotic procedures [1]. Although the review emphasized that the number of high-quality clinical studies remains limited, the findings indicate that the technology is beginning to move beyond laboratory demonstrations.

A related synthesis of 25 peer-reviewed studies published during 2024 and 2025 reported approximately 20–25% reductions in operative time, reductions in intraoperative complications of around 20–30%, and improvements in targeting accuracy in selected AI-assisted robotic procedures compared with conventional or non-AI-assisted approaches [5]. These findings should be interpreted cautiously because differences in study design, tumour type, surgical procedure, patient population, and outcome definitions make direct comparison difficult. Nevertheless, the consistency of improvement across several reported outcomes provides evidence that AI-assisted robotics may offer practical advantages beyond the mechanical benefits of robotic surgery alone.

Thoracic oncology provides a further example of this integration. Studies combining AI, augmented reality, and robotic assistance have reported improvements in tumour localization and resection accuracy together with reductions in operative time and postoperative complications [3]. In such systems, preoperative and intraoperative information can be transformed into visual or spatial guidance for the surgeon. The importance of this approach is that tumour location can change relative to the original preoperative image because of respiration, tissue deformation, surgical manipulation, or changes in patient positioning. Therefore, the ability to update the target during the procedure is an important requirement for a genuine closed-loop system.

The confirmation stage provides some of the clearest evidence of both the promise and the limitations of the current technology. In robotic surgery, confirmation can involve intraoperative imaging, augmented-reality overlays, tissue assessment, margin evaluation, or other sensing mechanisms that determine whether the planned intervention corresponds to the actual anatomical situation [1,3]. AI can compare the observed intraoperative state with the original tumour map and identify discrepancies. This capability is particularly important during cancer resection because a procedure may appear technically successful while microscopic or spatially difficult-to-detect tumour tissue remains at the surgical margin. A confirmation mechanism therefore changes the system from one that simply executes a preoperative plan into one

that can potentially detect and respond to deviations during treatment.

At the micro- and nanoscale, the evidence presents a different picture. Nanorobotic systems have demonstrated promising capabilities for targeted drug delivery, tumour recognition, and controlled therapeutic release, but most of the available evidence remains preclinical [7–9]. DNA-origami-based nanorobots carrying therapeutic payloads have been investigated in animal tumour models, where tumour growth reduction has been demonstrated while attempting to limit exposure of surrounding healthy tissue [8]. These systems use biological characteristics of the tumour microenvironment to regulate activation, providing an example of confirmation occurring at the molecular rather than macroscopic imaging level.

Magnetically controlled micro- and nanorobotic systems represent another important research direction. External magnetic fields can potentially provide directional control over small carriers, allowing them to move toward a desired anatomical region rather than depending entirely on passive drug diffusion [7,9]. Biohybrid systems that combine biological propulsion mechanisms with drug-loaded carriers have also been investigated for targeted cancer therapy [9]. These approaches are particularly relevant to the targeting stage because they introduce a physical mechanism for moving therapeutic agents toward localized disease sites.

The potential improvement over passive delivery is significant. Conventional systemic drug administration can expose healthy tissues to therapeutic agents because the drug is distributed throughout the body rather than being concentrated exclusively at the tumour. Reviews of targeted nanomedicine have reported substantially improved tumour accumulation for actively guided or intelligently designed carriers compared with passive delivery approaches [7,8]. However, these reported improvements are highly dependent on the experimental model, carrier design, targeting mechanism, tumour type, and measurement method. Therefore, the results should be viewed as evidence of technological potential rather than as established clinical effectiveness.

Artificial intelligence and reinforcement learning introduce another promising dimension to nanorobotic targeting. Simulation studies have demonstrated that reinforcement-learning-based nanorobot controllers can modify a carrier's trajectory in response to environmental information, including local biomarker concentrations [6,11]. Instead of following a completely predetermined route, the agent can use feedback to select subsequent actions. Conceptually, this represents a nanoscale equivalent of the closed-loop control used in larger robotic systems. The nanorobot senses its environment, interprets the available information, modifies its movement, and continues sensing as it approaches the target. Although these studies remain largely simulation-based, they provide an important proof of concept for integrating AI-based decision-making with microscale robotic actuation.

The confirmation mechanism in nanorobotic systems is also fundamentally different from that used in surgical robotics. A surgical robot may use imaging, optical sensing, ultrasound, or other intraoperative information to verify the position of an instrument or the status of tissue. A nanorobotic system may instead rely on biochemical cues. For example, tumour-associated characteristics such as local acidity or specific molecular biomarkers can be used as triggers for therapeutic activation [7,8]. In this situation, confirmation is not necessarily a visual determination that the robot has reached a particular coordinate. Instead, the biochemical environment itself provides evidence that the carrier is located within the intended tumour microenvironment.

This difference illustrates why the detect–target–confirm framework must remain scale-independent. At the macro scale, the loop may involve medical images, robotic surgical instruments, and intraoperative imaging. At the micro/nanoscale, the same conceptual loop may involve biomarker detection, magnetic or biological propulsion, and molecularly triggered drug release. Although the hardware and sensing mechanisms differ substantially, both systems follow the same fundamental sequence: identify the target, move or act toward the target, and use feedback to determine whether the intended location or biological condition has been reached.

Another important finding from recent literature is that the three stages are not equally mature. Detection is currently the most developed component because it benefits from large public datasets, standardized imaging benchmarks, and rapidly advancing deep-learning architectures [12–18]. Robotic targeting is also relatively mature at the surgical scale because robotic platforms are already used clinically, although AI-controlled or AI-guided autonomy remains limited [1,5]. Confirmation is less standardized because there is no single sensing modality that is appropriate for all tumour types, anatomical locations, and treatment strategies. The nanorobotic component is less mature still, with many promising results coming from animal studies, laboratory experiments, and simulations rather than controlled human clinical trials [6–9].

The distinction between technical performance and clinical effectiveness is therefore critical. High Dice scores

or classification accuracy demonstrate that an AI model can perform well on a particular dataset, but they do not automatically establish that the model will generalize to new hospitals, scanners, patient populations, or unusual tumour presentations. Similarly, improved robotic targeting accuracy in a controlled procedure does not necessarily demonstrate improved long-term cancer survival. For the proposed framework, evidence must ultimately be evaluated at several levels: detection accuracy, spatial targeting accuracy, intervention safety, confirmation reliability, clinical outcomes, and the ability of the complete loop to adapt when conditions change.

At present, very few studies evaluate these components together as a single end-to-end system. Most studies report performance independently: AI researchers report segmentation or classification metrics, robotic studies report surgical precision or operative outcomes, and nanorobotic studies report drug-delivery efficiency, tumour reduction, or navigation performance. This fragmentation makes direct evaluation of the complete detect–target–confirm architecture difficult. A system could theoretically achieve excellent detection but poor targeting, or highly accurate targeting without reliable confirmation. Therefore, the performance of the complete loop cannot be inferred simply by adding the performance of its individual components.

A useful way to conceptualize the current evidence is therefore as a technology-readiness gradient. The detection stage has moved toward advanced clinical decision support and increasingly multimodal tumour characterization. Robot-assisted targeting has progressed into real clinical workflows, although substantial human involvement remains necessary. Confirmation technologies are emerging but lack universal standards for real-time verification. Micro/nanorobotic targeting and AI-guided nanoscale navigation remain predominantly experimental. The field is consequently not yet at the stage of a fully autonomous cancer-treatment loop, but the evidence indicates that each individual component required for such a system is progressing.

Taken together, the surgical-scale evidence shows a technology that is already producing measurable clinical benefits in a growing number of real-patient studies, while the nanoscale evidence shows a technology whose underlying detect–target–confirm logic can be demonstrated convincingly in laboratory, animal, and simulation environments but has not yet been sufficiently validated in humans [1,6–9]. The strongest current evidence therefore supports a hybrid interpretation: AI-based tumour detection and characterization are increasingly mature; robotic intervention is clinically established in selected surgical applications; and intelligent micro/nanorobotic therapy remains an emerging research technology.

The major research opportunity is to move from these individually validated components toward integrated evaluation. Future studies should not only ask whether an AI model can detect a tumour or whether a robot can reach a target, but whether the complete system can reliably detect a clinically meaningful target, translate that target into safe physical action, verify the result, and adapt when the observed outcome differs from the intended one. Such end-to-end evaluation would provide a more meaningful measure of the maturity of the detect–target–confirm paradigm and would help identify which component represents the limiting factor in a particular clinical application.

Overall, the recent evidence supports the central premise of this paper: the individual technologies required for AI-enabled closed-loop precision oncology are no longer purely speculative. However, their levels of validation remain uneven, and the transition from promising components to a clinically reliable integrated system will require stronger prospective clinical evidence, standardized evaluation metrics, robust real-time confirmation, and carefully controlled integration between AI decision-making and robotic action.

4. Challenges and Limitations

Several obstacles stand between the current state of the art and a fully autonomous, clinically deployed detect-target-confirm system.The first is generalisability. Imaging AI models trained on data from a limited number of institutions, scanner manufacturers, or patient populations often lose accuracy when applied elsewhere, which is precisely why recent foundation-model efforts have deliberately trained on data pooled from many cohorts and why federated-learning approaches, which let models train across multiple hospitals without centralising patient data, have become an active area of research for tumour segmentation specifically [3, 24]. A detection model that performs unevenly across populations creates a downstream risk: a robotic system acting on a flawed spatial map is not a safety-neutral error, it is a physical one. The second is the physical difficulty of navigation, which differs by scale but does not disappear at either end of it. Surgical robots must operate in soft, deformable tissue that shifts during a procedure in ways that are hard to fully predict from preoperative imaging alone [1]. Nanoscale carriers face an even harder version of the same problem: they must navigate a fluid, turbulent, and biologically

hostile environment inside blood vessels, and fabrication and control challenges mean that reliably distinguishing healthy cells from tumour cells at that scale, and avoiding failures in payload dosage, remains an unresolved engineering problem rather than a solved one [7].

The third is the fragmentation of the evidence base itself. Much of the strongest robotic-surgical evidence to date comes from a relatively small number of systematic reviews built on real patient data, while the strongest nanorobotic evidence is preclinical [1,8]. Very few published studies attempt to evaluate the detect-target-confirm loop as a single, end-to-end system with a unified performance metric; instead, detection accuracy, robotic targeting accuracy, and confirmation or verification accuracy tend to be reported separately, if at all, which makes it difficult to state with confidence how well the loop performs as a whole rather than as three disconnected parts.

Finally, there are regulatory, safety, and workflow questions that are not purely technical. Increasing the autonomy of a surgical or nanorobotic system raises questions about liability when the AI-generated map or decision is wrong, about how much human oversight is required at each stage, and about the cost and training burden of integrating these systems into existing hospital workflows [1,5]. None of these issues are unique to oncology, but the stakes of an error, whether a missed margin during surgery or an off-target nanorobotic drug release, make them especially pressing in this domain.

5. Future Directions

Several developments already visible in the field point toward the next generation of AI-enabled robotic systems for precision oncology. A major direction is the development of foundation-style AI models capable of learning from large, diverse, and multimodal medical datasets. Such models could improve the generalisability of tumour detection and characterization across different hospitals, imaging devices, patient populations, and cancer types. Instead of relying on a single imaging modality, future systems are likely to combine radiological images, histopathology, molecular biomarkers, clinical information, and intraoperative observations to create a more comprehensive representation of the tumour and its surrounding anatomical environment.

Another important direction is the development of adaptive robotic systems capable of responding to changes during an intervention. Current robotic procedures often depend on preoperative planning and continuous human supervision. Future systems could increasingly incorporate real-time information from imaging and sensors to update the planned trajectory, compensate for tissue movement and deformation, and modify the intervention according to the observed condition of the patient. This would allow robotic systems to move from primarily executing predefined instructions toward more adaptive and context-aware physical intervention.

The confirmation stage represents one of the most important opportunities for future development. Reliable real-time confirmation mechanisms are required to determine whether the intended tumour region has been reached, whether the appropriate tissue has been treated, and whether residual malignant tissue remains. Advances in intraoperative imaging, optical sensing, ultrasound, biochemical sensing, and molecular detection could provide the feedback necessary to make this stage more reliable. For micro- and nanorobotic systems, improvements in biomarker recognition and controlled therapeutic activation could similarly provide stronger evidence that a therapeutic payload has been released at the intended location.

Future research should also focus on developing a common performance framework for the complete detect–target–confirm loop. At present, detection accuracy, targeting accuracy, intervention outcomes, and confirmation performance are frequently evaluated as separate components. A standardized end-to-end evaluation framework would make it possible to determine whether an integrated system can successfully identify a target, translate the target into physical action, verify the outcome, and adapt when the observed result differs from the intended result. Metrics should consider not only technical accuracy but also safety, robustness, response time, reliability, clinical effectiveness, and the degree of human supervision required.

The integration of macro-scale robotic surgery with micro- and nanoscale therapeutic systems is another promising long-term direction. A future clinical workflow could potentially connect different levels of intervention within the same treatment pathway. For example, a robotic system could obtain a precisely localized biopsy, AI could analyse the resulting information to identify relevant tumour characteristics, and a targeted therapeutic carrier could subsequently be directed toward the identified biological target. Such an integrated approach could connect diagnosis, treatment planning, physical intervention, and therapeutic verification within a single coordinated framework.

Personalized and adaptive oncology is also likely to benefit from this convergence. Tumours can change over time and may respond differently to treatment across

different regions of the same lesion. Future AI systems could continuously update tumour maps using newly acquired imaging, pathology, and biomarker information. Robotic systems could then use these updated maps to modify treatment strategies rather than relying on a fixed plan generated before the procedure. This could support more localized and patient-specific treatment while reducing unnecessary exposure of healthy tissue.

Another important research direction is the development of human-supervised autonomy. Fully autonomous intervention may not be appropriate for many clinical situations, particularly when uncertainty or unexpected anatomical conditions are encountered. Future systems could therefore adopt different levels of autonomy depending on the complexity and risk of the task. Routine and highly predictable actions could be automated, while uncertain or high-risk decisions could require explicit clinician approval. Such an approach would allow AI and robotics to provide increasing levels of assistance without removing clinical responsibility from qualified healthcare professionals.

Safety, explainability, cybersecurity, and regulatory validation will remain essential as these systems become more integrated. Future platforms will need mechanisms for detecting uncertainty, identifying potentially incorrect AI predictions, preventing unsafe robotic movements, and allowing clinicians to intervene immediately when required. Transparent reporting of system performance across diverse patient populations will also be necessary to establish confidence in clinical deployment. In addition, the development of standardized testing environments and clinically meaningful benchmarks will help researchers compare different AI–robotic architectures more consistently.

Ultimately, the long-term objective is to transform the current collection of AI models, imaging technologies, robotic platforms, and targeted delivery systems into a coordinated closed-loop clinical architecture. Such a system would continuously sense the patient, interpret the available information, select an appropriate action, execute that action with high precision, and verify the resulting outcome. The transition toward this vision is likely to occur incrementally, beginning with AI-assisted clinical decision support and robot-assisted procedures and gradually progressing toward increasingly adaptive and integrated systems.

The future of precision oncology therefore lies not simply in developing more accurate AI models or more precise robots independently, but in establishing reliable communication and feedback between perception, decision-making, physical intervention, and outcome verification. The detect–target–confirm framework provides a conceptual foundation for this integration and highlights the importance of treating detection, targeting, and confirmation as interconnected stages of a continuous clinical process.

6. Conclusion

Artificial intelligence and robotics are converging into a workflow that can reasonably be described as a detect-target-confirm loop for localized cancer care, with AI serving as the system's interpretive layer and robotics, at scales ranging from full-sized surgical arms to nanometre-scale carriers, serving as its means of physical action. The detection stage is the most mature, supported by a large and fast-growing body of deep-learning research on tumour segmentation and multimodal diagnosis. The targeting stage is maturing quickly on the surgical side, where real-patient evidence already shows measurable gains in precision and outcomes, and is progressing rapidly but still preclinically on the nanorobotic side. The confirmation stage, which is what actually closes the loop and allows adaptive correction rather than a one-shot plan, remains the least standardised of the three and is the area most likely to determine how quickly this technology moves from an assembly of promising components into a single, reliably closed-loop clinical

system. None of the individual pieces reviewed here are speculative; what remains speculative is how soon, and how completely, they will be integrated.

DECLARATIONS

- Funding Source: No funding received.
- Acknowledgement: The authors acknowledge the support and academic facilities .
- Data Availability: The study is a review of published literature; the supporting information is available through the cited references and from the corresponding author upon reasonable request.
- Authors' Contribution: **Dr. Chaitanya Kumar M V** contributed to conceptualization and supervision. **Dr. Ashoka S. B** contributed to literature review and analysis. **Dr. Niranjan R Chougala** contributed to technical review and critical editing. **Mr. Shashank H S** contributed to literature collection, manuscript drafting, formatting, and reference management.
- Use of AI and AI-Assisted Technologies: AI-assisted tools were used to support language editing, formatting, and manuscript organization. The authors

reviewed and approved the final scientific content.

- Conflict of Interest: The authors declare no conflict of interest.
- Copyright Permissions: Copyright permissions for all figures and tables included in this article have been duly obtained, where applicable.

References

[1] Leszczyńska A, Obuchowicz R, Strzelecki M, Seweryn M. The integration of artificial intelligence into robotic cancer surgery: a systematic review. Journal of Clinical Medicine. 2025 Sep 1;14(17):6181. https://www.mdpi.com/2077-0383/14/17/6181

[2] Najar Najafi N, Hajihassani H, Azimzadeh Irani M. The impact of artificial intelligence on cancer diagnosis and treatment: a review. Cancer informatics. 2025 Sep;24:11769351251371273. https://journals.sagepub.com/doi/full/10.1177/11769351251371273

[3] Tiwari A, Mishra S, Kuo TR. Current AI technologies in cancer diagnostics and treatment. Molecular cancer. 2025 Jun 2;24(1):159. https://link.springer.com/article/10.1186/s12943-025-02369-9

[4] Ng J, Wah K. AI-enhanced robotic hands: a breakthrough in early tumour detection and removal. Journal of Robotic Surgery. 2026 Feb 5;20(1):218. https://link.springer.com/article/10.1007/s11701-026-03159-1

[5] Wah JN. The rise of robotics and AI-assisted surgery in modern healthcare. Journal of Robotic Surgery. 2025 Jun 20;19(1):311. https://link.springer.com/article/10.1007/s11701-025-02485-0

[6] Kavousinejad S. Simulation of nanorobots with artificial intelligence and reinforcement learning for advanced cancer cell detection and tracking. arXiv preprint arXiv:2411.02345. 2024 Nov 4. https://arxiv.org/abs/2411.02345

[7] Das KP. Nanoparticles and convergence of artificial intelligence for targeted drug delivery for cancer therapy: Current progress and challenges. Frontiers in Medical Technology. 2023 Jan 6;4:1067144. http://frontiersin.org/journals/medical-technology/articles/10.3389/fmedt.2022.1067144/full

[8] Anto M, Mukherjee K, Mukherjee K. Nano bio-robots: a new frontier in targeted therapeutic delivery. Frontiers in Robotics and AI. 2025 Nov 6;12:1639445. https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2025.1639445/full

[9] Leong WY. Biohybrid micro-robots for targeted drug delivery in cancer therapy. Engineering Proceedings. 2025 Dec 24;120(1):4. https://www.mdpi.com/2673-4591/120/1/4

[10] Olawade DB, Clement David-Olawade A, Adereni T, Egbon E, Teke J, Boussios S. Integrating AI into cancer immunotherapy—A narrative review of current applications and future directions. Diseases. 2025 Jan 20;13(1):24. https://www.mdpi.com/2079-9721/13/1/24

[11] Abdul Rasool Hassan B, Mohammed AH, Hallit S, Malaeb D, Hosseini H. Exploring the role of artificial intelligence in chemotherapy development, cancer diagnosis, and treatment: present achievements and future outlook. Frontiers in oncology. 2025 Feb 4;15:1475893. https://www.frontiersin.org/journals/oncology/articles/10.3389/fonc.2025.1475893/full

[12] Lu NH, Huang YH, Liu KY, Chen TB. Deep learning-driven brain tumor classification and segmentation using non-contrast MRI. Scientific reports. 2025 Jul 30;15(1):27831. https://www.nature.com/articles/s41598-025-13591-2

[13] Hamza A, Damaševičius R. Deep learning for brain tumor segmentation and classification: a systematic review of methods and trends. Computers, materials and continua. 2026;86(1):1-41. https://epubl.ktu.edu/object/elaba:265482718/

[14] Bouhafra S, El Bahi H. Deep learning approaches for brain tumor detection and classification using MRI images (2020 to 2024): a systematic review. Journal of Imaging Informatics in Medicine. 2025 Jun;38(3):1403-33. https://link.springer.com/article/10.1007/s10278-024-01283-8

[15] Wahid KA, Dede C, El-Habashy DM, Kamel S, Rooney MK, Khamis Y, Abdelaal MR, Ahmed S, Corrigan KL, Chang E, Dudzinski SO. Overview of the head and neck tumor segmentation for magnetic resonance guided applications (HNTS-MRG) 2024 challenge. InChallenge on Head and neck tumor segmentation for MRI-guided applications 2024 Oct 9 (pp. 1-35). Cham: Springer Nature Switzerland. https://link.springer.com/chapter/10.1007/978-3-031-83274-1_1

[16] Karam L, Wang Y, Kasivisvanathan V, Rusu M, Hu Y, Saeed SU. Promptable cancer segmentation using minimal expert-curated data. InAnnual Conference on Medical Image Understanding and Analysis 2025 Jul 15 (pp. 44-58). Cham: Springer Nature Switzerland. https://link.springer.com/chapter/10.1007/978-3-031-

98694-9_4

[17] Aftab M, Mehmood F, Zhang C, Nadeem A, Dong Z, Jiang Y, Liu K. AI in oncology: transforming cancer detection through machine learning and deep learning applications. arXiv preprint arXiv:2501.15489. 2025 Jan 26. https://arxiv.org/abs/2501.15489

[18] Dao L, Ly NQ. Recent advances in medical image classification. arXiv preprint arXiv:2506.04129. 2025 Jun 4. https://arxiv.org/abs/2506.04129

[19]

[20] Nanobots in healthcare research report 2024: a $7.1 billion market in 2023 — global forecast to 2030. GlobeNewswire. 2024. https://finance.yahoo.com/news/nanobots-healthcare-research-report-2024-104800388.html?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLmNvbS8&guce_referrer_sig=AQAAAIjJ7rR_qHOWOPdK7WJzO4-mOmoHFy1u-0mlqOKtwkVoYdo0apTN9QewgYOEmxYKFJzyxPsI-0952Xdq7XAd_Ixg-ECFBCUezLakQnqSbWOR-ZAZoDrjOMYhs2MsGztpK3CCpnIRn6N5Ti5UUs3qQBNbpI0o1R3qbndgmeXPn83W

[21] Rangnekar A, Miranda J, Horvat N, Chahwan S, Alrayess S, Apte A, Iyer A, LoCastro E, Ravella R, Gollub MJ, Petkovska I. Tumor-aware augmentation with task-guided attention analysis improves rectal cancer segmentation from magnetic resonance images. Biomedical Signal Processing and Control. 2026 Dec 1;128:111229. https://www.sciencedirect.com/science/article/abs/pii/S1746809426017830

[22] Linardos A, Pati S, Baid U, Edwards B, Foley P, Ta K, Chung V, Sheller M, Khan MI, Jafaritadi M, Kontio E. The MICCAI Federated Tumor Segmentation (FeTS) Challenge 2024: Efficient and Robust Aggregation Methods for Federated Learning. arXiv preprint arXiv:2512.06206. 2025 Dec 5. https://arxiv.org/abs/2512.06206