Imagine an oncologist stepping out of tumor board and into clinic with a complex case at hand. The patient’s tumor has multiple high-risk features, genomic mutations, and borderline indications for therapy. In the past, the doctor might sift through guidelines, trial data, and pathology reports alone – but today he also consults an AI “colleague” for backup. This AI doesn’t make the decision, but it offers evidence-based insights gleaned from vast data, helping the clinician consider critical details that might otherwise be missed.

From Hype to Help: AI as the Oncologist’s Assistant

Early experiments with AI in oncology – like IBM’s Watson for Oncology – generated buzz by promising to recommend treatments from big data. In practice, however, many such systems struggled to gain clinician trust or improve outcomes. They often relied on curated guidelines and literature without integrating the full picture of an individual patient. The lesson learned was that AI works best as a consultant, not a replacement, and only when grounded in validated clinical evidence.

In fact, the early mainstream wave of AI, driven by generative models, reinforced this point: these systems can draft convincing answers to medical questions, while also highlighting the need for guardrails—so responses stay accurate, contextual, and evidence-based.

– **ASCO (Guidelines Assistant on Vertex AI/Gemini)**

Recognizing the need for reliability, oncology leaders have begun harnessing such AI within strict evidence-based confines. For example, the American Society of Clinical Oncology (ASCO) partnered with Google Cloud to develop a Guidelines Assistant on Vertex AI (Gemini) that lets clinicians query ASCO’s guideline library and receive instant, citation-linked answers drawn only from vetted content. It delivers the convenience of an assistant without the usual noise—evidence at your fingertips, not a free-form recommender.

– **OpenAI (Penda Health EHR-Embedded AI Consult)**

Outside oncology, a real-world, EHR-embedded “AI consult” from Penda Health and OpenAI showed what good looks like: a simple traffic-light interface (green/yellow/red) that fires only at decision points and preserves clinician autonomy. In a pragmatic evaluation across 39,849 visits in 15 clinics, clinicians using the tool made 16% fewer diagnostic errors and 13% fewer treatment errors than controls—evidence that workflow-native decision support can reduce real clinical errors. For oncology, the value lies in the method: tight triggers, minimal cognitive load, a clear rationale, and an active rollout.

Building on these lessons, today’s emerging AI tools take a more pragmatic approach. Rather than offering generic advice, they serve up specific, clinically validated guidance at key decision points—much like a trusted colleague rounding with the team. Critically, the most promising AI decision-support tools are those built and tested with real patient data and aligned with established protocols.

Oncologists are understandably cautious—any AI suggestion must be anchored in something tangible, like a published trial or a guideline. This is why recent “AI consult” systems focus on narrow but impactful tasks: predicting therapy benefit, flagging high-risk features, or aggregating patient data for review. These are not sci-fi algorithms operating in a vacuum; they are sophisticated extensions of the tumor board, able to synthesize pathology, imaging, genomic, and clinical information into actionable insights. Below, we explore a few such examples currently shaping oncology practice. Among them, ArteraAI is guideline-level validated, while the others remain research- or early-clinical-stage with promising evidence.

Multimodal AI Tools Coming of Age

— ArteraAI Prostate Test

One striking example is the ArteraAI Prostate Test – an AI-driven tool that has already made its way into clinical guidelines. ArteraAI's system analyzes a prostate cancer patient's digitized biopsy histology alongside clinical variables to predict how the patient will fare with different treatment options. In effect, it produces a personalized risk report: will this man benefit from adding short-term hormone (androgen deprivation) therapy to radiation, or can he safely opt out of the extra treatment?

This AI model was trained and validated on thousands of patients from large randomized phase III trials, with long-term follow-up, so its predictions aren't abstract probabilities – they're grounded in real outcomes data. The core model is multimodal: it combines pathology images with clinical data. In published evaluations, this approach outperformed conventional NCCN-style risk grouping in predicting long-term prostate cancer outcomes, showing roughly a 9–15% relative improvement in discriminatory performance (i.e., the ability to separate higher-risk from lower-risk patients for endpoints like progression and metastasis) compared with traditional clinical risk tools. In other words, the AI proved more adept than standard clinical criteria at forecasting which patients are likely to relapse or die of their cancer over time.

Critically, ArteraAI's algorithm also identified a biologically and clinically meaningful subgroup: intermediate-risk prostate cancer patients who truly benefit from adding short-term androgen deprivation therapy (ADT) to radiotherapy, versus those who gain little from the added hormones. In patients predicted to benefit, intensification with ADT improves long-term control; in patients predicted not to benefit, the extra hormones – and their side effects – can potentially be avoided.

This level of evidence was strong enough that in 2024 the National Comprehensive Cancer Network (NCCN) cited the ArteraAI Prostate Test in its Prostate Cancer Guidelines as a prognostic and predictive adjunct for localized disease. That listing is widely described as the first AI-enabled biomarker of its kind to be incorporated into NCCN prostate oncology guidance. For clinicians and patients, this offers a new level of confidence: an AI consult, based on the patient's own tumor pathology and clinical profile, helps stratify risk and personalize therapy choice. The doctor still makes the call, but now with an AI-derived analysis of who is most likely to benefit from treatment intensification – and who might safely avoid unnecessary toxicity.

Predictors (Research-Stage): H&E Slide–Based Models in NSCLC Immunotherapy

On the lung cancer front, similar experimental AI consults are tackling one of the toughest questions: which patients will respond to immunotherapy. Checkpoint inhibitors (PD-1/PD-L1 blockers) have transformed the treatment of advanced non–small cell lung cancer (NSCLC), but only about 20–30% of patients experience durable benefit. Oncologists today lean on imperfect biomarkers – PD-L1 expression levels, tumor mutational burden (TMB), and similar measures – to guess who might respond.

Multimodal AI models are now being developed to improve this decision point by learning from diverse patient data. For example, researchers have shown that deep learning algorithms can analyze routine diagnostic material – the same H&E pathology slides already produced for standard-of-care biopsy – to detect hidden morphologic and microenvironmental patterns predictive of immunotherapy response. In a recent multicenter study spanning several hospitals, a deep learning model that extracted features from H&E tumor specimens emerged as an independent predictor of response to PD-1/PD-L1 immunotherapy and of progression-free survival, even after adjusting for PD-L1 status, TMB, and other known covariates. In practical terms, this means that a digitized biopsy, when processed by AI, might reveal whether a patient's tumor looks "immune-responsive" or "immune-cold" in ways that are not obvious to the human eye.

If prospectively validated in interventional trials, such a tool could flag patients unlikely to respond to checkpoint inhibitors before they embark on months of treatment – a valuable "second opinion" to help decide whether to proceed with expensive, immune-based therapy or pivot earlier to an alternative strategy. But at present this remains investigational: these H&E-based immunotherapy predictors have strong retrospective and external validation data, yet they have not been adopted into major NSCLC guidelines, and clinicians are not (yet) using them to deny or escalate therapy on their own.

— NSCLC Immunotherapy: CT Radiomics (QVT) (Research-Stage)

Other AI efforts in NSCLC are combining data from medical imaging and clinical labs to further refine immunotherapy decisions. One example is AI-based radiomic analysis of standard CT scans. These approaches quantify tumor characteristics that are effectively invisible to the naked eye – such as the tortuosity and "chaoticness" of blood vessels feeding the tumor – and link them to immunotherapy outcomes.

In one notable 2023 study and related multicenter reports, investigators described a CT imaging biomarker called quantitative vessel tortuosity (QVT). Tumors with highly tortuous, disorganized vasculature were more likely to be non-responders to PD-1/PD-L1 checkpoint inhibitors and to have shorter survival, even after accounting for PD-L1 levels and other clinical factors. In other words, the vascular "fingerprint" on a baseline CT scan carried a predictive signal about who would and would not benefit from immunotherapy.

If prospectively validated, these kinds of imaging-derived predictors – taken together with clinical context – move us closer to an AI that can say, "Given this patient's scan and profile, immunotherapy has a low chance of success – consider an alternative or an intensified approach." It's important to emphasize that these radiomic and multimodal predictors are still research- or early-clinical-stage. None of them are replacing established biomarkers like PD-L1 in current practice. But they illustrate what an "AI consult" could soon look like: a synthesis of imaging, pathology, and genomic data to support a yes/no immunotherapy decision and to better stratify patients for clinical trials.

The Integrated Oncology "Copilot" at the Point of Care

Perhaps the most ambitious use of AI in clinical decision support is appearing at the hospital or clinic level, where AI acts as a continuously updated data synthesizer for every patient case. Consider the approach taken by Yonsei Cancer Center in South Korea: they developed an in-house, AI-enabled clinical decision support system that continuously pulls each patient's pathology reports, radiology images, genomic test results, prior treatments, and clinical history into one unified dashboard. This platform – described as the Yonsei Cancer Data Library – is not built for a single tumor type or a single decision point. It is intended to support oncology care broadly across the center.

The system aggregates more than 800 structured data elements per patient and refreshes in near-real time as new results come in. When an oncologist opens this dashboard, they see a longitudinal timeline of the patient's cancer journey: key lab trends, imaging milestones, molecular markers, prior lines of therapy, and outcomes – all organized and visualized by AI. The interface can surface potential red flags, such as a slow drift away from guideline-concordant care, or a concerning pattern (like steadily rising tumor markers) that might warrant intervention sooner rather than later.

In Yonsei's initial experience, oncology staff reported high satisfaction (scores above 4 out of 5) with this integrated AI-assisted workflow. By letting AI quietly manage the data deluge in the background, clinicians reported being freer to focus on interpreting the insights and talking with patients, rather than clicking through scattered PDFs and siloed record systems. This kind of multimodal clinical decision support system is essentially an AI copilot for the entire oncology team. It doesn't diagnose or decide, but it helps ensure that no critical piece of pathology, imaging, genomics, labs, or symptom history is lost in the noise when building a treatment plan.

In practice, systems like this can rapidly retrieve how a patient's tumor genomic profile maps to available targeted therapies or open clinical trials, while simultaneously reminding the clinician of past toxicity issues documented in the chart. Some centers are now piloting in-silico forecasting on top of this infrastructure – for example, generating individualized survival curves or risk projections under different standard-of-care regimens, trained on large institutional outcome datasets. That kind of projection can support difficult conversations about prognosis and treatment intensity.

Crucially, these copilots are designed to stay anchored to clinical guidelines and institutional protocols. The AI is not inventing experimental treatments; it is matching patient-specific data to evidence-based options and flagging when care may be drifting from best practice. For oncologists who are juggling increasingly complex cases and ever-expanding data streams, that kind of context at the point of care is becoming less "nice to have" and more essential.

Conclusion: A Future of Informed Decision-Making, Not Autopilot

Across these examples runs a common theme: AI in oncology decision support works best when it augments the clinician's judgment with multi-dimensional analysis, rather than trying to automate the decision itself. Whether it's a specialized tool that interprets a prostate biopsy to estimate who truly benefits from adding short-term hormone therapy to radiation, a research-stage model that stratifies which lung cancer patients are likely (or unlikely) to respond to PD-1/PD-L1 immunotherapy, or a hospital-wide platform that unifies pathology, imaging, genomics, labs, and prior treatments into a single continuously updated dashboard, the aim is the same. These AI consults act as intelligent advisors, grounded in real clinical data, that strengthen the clinician's hand at the moment of decision. They surface things a busy human might miss: subtle morphology on an H&E slide, complex risk curves derived from thousands of trial patients, or a quiet reminder that today's plan is drifting from evidence-based best practice.

For oncologists, the promise of these multimodal AI systems is more confidence and clarity in choosing the right treatment for the right patient at the right time. Some of the most credible tools have undergone peer-reviewed validation and, in certain cases, have now been incorporated into major guidelines — for example, an AI-enabled prostate cancer test (ArteraAI) that appears in the NCCN Prostate Cancer Guidelines to help distinguish which localized patients truly need short-term androgen deprivation therapy with radiotherapy and which might safely avoid it. Others, like deep learning–based immunotherapy response predictors in advanced non–small cell lung cancer or CT radiomics biomarkers such as quantitative vessel tortuosity, are still in prospective evaluation and are not yet standard of care, but they are already outperforming classic single biomarkers (like PD-L1 alone) in multicenter validation studies. No doubt challenges remain — integrating AI into workflow, training clinicians, guarding against algorithmic bias, and ensuring interpretability and auditability — but the trajectory is set.

In the clinic of tomorrow, an oncologist facing a high-stakes decision will not be forced to choose between skimming dozens of siloed PDFs or relying on intuition alone. Instead, with a single “AI consult,” they will be able to pull forward distilled evidence from millions of data points and prior cases, summarized in a clinically meaningful way. The final judgment will still rest with the human oncologist – but it will be made with a clearer view of risk, benefit, and precedent.

In sum, the oncology AI consult is moving from futuristic concept to practical reality. By embracing multimodal data and focusing on validated, guideline-aware algorithms where they exist, these tools are beginning to deliver on the long-promised vision of precision support in cancer care. They are not here to replace the art of oncology. They are here to sharpen it – to make sure that when an oncologist sits with a patient and says, “Here’s what I recommend, and here’s why,” that answer reflects not just experience, but the best available evidence, assembled instantly and tailored to that one person. As these systems continue to mature, oncologists can look forward to making treatment decisions with greater insight and assurance, knowing that no matter how fast the field grows, they have an ever-ready digital ally to help navigate the complexity.

Author’s Note

Dr. Amil Družić

The momentum in oncology AI has become impossible to ignore. In November 2025, the European Society for Medical Oncology will host its first standalone ESMO AI & Digital Oncology Congress—a dedicated forum for exploring how artificial intelligence and digital tools are reshaping cancer care. Its very existence underscores how rapidly this field is moving into mainstream oncology.

As a physician specialising in oncology and radiotherapy with a deep interest in technology, I have followed these developments closely. In partnership with the Association of Oncologists in Bosnia and Herzegovina, we conducted a national survey examining how—and in what ways—oncologists use digital and AI tools in everyday practice, spanning research, clinical decision support, scientific writing, communication, and public awareness. I will present these findings at the ESMO AI & Digital Oncology Congress.

From my perspective, it's an incredibly exciting time to be at the intersection of oncology and digital innovation. We are witnessing a convergence of need and opportunity: clinicians overwhelmed by data and options, and technology that's finally capable of meaningfully assisting with that burden. My background in clinical oncology and research has shown me the value of evidence-based decision-making, and AI, when applied responsibly, is poised to enhance that process. The key will be to ensure these tools are developed with oncologists, not just for them – aligning with real-world workflow and high standards of clinical evidence.

I am optimistic that with continued collaboration between clinicians, researchers, and tech experts, we will navigate the challenges (data quality, bias, integration) and unlock AI's full potential in cancer care. The discussions at forums like the ESMO AI & Digital Oncology Congress, and the feedback from front-line oncologists in surveys and studies, all point to a common goal: using AI to make cancer treatment smarter, more personalised, and more efficient—without losing the human touch that defines medicine. In the end, the future of oncology will not be about AI autopilot but about informed decision-making, with AI as a powerful ally to help us help our patients better. Put simply: AI will not replace oncologists—but it will redefine their capabilities, enabling them to synthesize vast amounts of clinical data with unprecedented speed and precision.

References

1. ASCO Post – How the AI-Powered ASCO Guidelines Assistant Is Improving Clinical Decision-Making(https://ascopost.com/issues/september-25-2025/how-the-ai-powered-asco-guidelines-assistant-is-improving-clinical-decision-making/)
2. ASCO Post – ASCO and Google Cloud AI-Powered Tool (https://ascopost.com/issues/june-10-2025/asco-and-google-cloud-ai-powered-tool/)
3. OpenAI – Pioneering an AI clinical copilot with Penda Health (https://openai.com/index/ai-clinical-copilot-penda-health/)
4. Penda Health – AI-Enabled Care overview (https://www.pendahealth.com/ai-enabled-care)
5. OpenAI/Penda – Methods PDF (https://cdn.openai.com/pdf/a794887b-5a77-4207-bb62-e52c900463f1/penda_paper.pdf)
6. Targeted Oncology – NCCN recommends first AI prognostic tool in prostate cancer (https://www.targetedonc.com/view/nccn-recommends-first-ai-prognostic-tool-in-prostate-cancer)
7. Urology Times – ArteraAI test included in NCCN Guidelines (https://www.urologytimes.com/view/arteraai-prostate-cancer-test-included-in-nccn-guidelines)
8. ArteraAI – For Patients (test overview) (https://artera.ai/for-patients)
9. JAMA Oncology – Deep Learning Model for NSCLC ICI response (https://jamanetwork.com/journals/jamaoncology/fullarticle/2828578)
10. PubMed – NSCLC ICI response DL model (https://pubmed.ncbi.nlm.nih.gov/39724105/)
11. Winship/Emory – AI biomarker predicts ICI response in lung cancer (https://winshipcancer.emory.edu/newsroom/articles/2023/new-ai-based-biomarker-can-help-predict-immunotherapy-response-for-patients-with-lung-cancer.php)
12. ASCO Post – News brief on QVT biomarker

(https://ascopost.com/news/january-2023/new-ai-based-biomarker-may-help-predict-immunotherapy-response-for-patients-with-nsclc/)
13. Stony Brook – Summary of Science Advances work (https://ai.stonybrook.edu/about-us/News/new-ai-based-biomarker-can-help-predict-immunotherapy-response-patients-lung-cancer)
14. npj Digital Medicine – Yonsei Cancer Data Library (Nature) (https://www.nature.com/articles/s41746-025-01508-2)
15. PubMed Central – Yonsei Cancer Data Library full text (https://pmc.ncbi.nlm.nih.gov/articles/PMC11868524/)
16. ESMO – AI & Digital Oncology Congress 2025 (https://www.esmo.org/meeting-calendar/esmo-ai-and-digital-oncology-congress-2025)
17. ESMO – Registration page (https://www.esmo.org/meeting-calendar/esmo-ai-and-digital-oncology-congress-2025/registration)