



# Clinical Artificial Intelligence

Past, Present and Future

---

**Irbaz Bin Riaz, MBBS, MS, MBI, PhD**

Director, Cancer Artificial Intelligence | Director, BEACON Lab  
Senior Associate Consultant, Hematology/Oncology & AI Informatics  
Mayo Clinic Arizona

Mayo UIUC Computational Genomics Course  
*June 26th, 2026*



# Disclosures

The speaker has the following relevant financial relationships to disclose:

**No relevant financial relationships to disclose**



# Today's Talk

*Big picture, three axes of impact, Some tensions to watch*



## The big picture — how clinical AI is evolving

From rule-based scores to predictive, generative, and agentic systems — and where uro-onc fits.



## Three axes of impact in Clinical AI

AI in the patient journey, in clinical workflows, and in the patient–clinician interaction.



## Some Practical Examples

The Modern AI Clinic: Agentic AI, AI based biomarkers(ADT Zero) and the ASCO living Guidelines



## Open problems & where you come in

The open tensions in the field — and the research problems they create.

# The Evolution of Clinical AI

*From AI that calculates to AI that acts — with the clinician as attending*



## Rule-Based AI

*“The Calculator”*

Fixed formulas. Same inputs always give the same score.

---

*D'Amico, NCCN risk groups, CAPRA, Briganti nomogram*



## Predictive AI

*“The Pattern Spotter”*

Learns from thousands of patients to predict what's coming next.

---

*Artera, Decipher, PSMA-PET response models, ctDNA classifiers*



## Generative AI

*“The Medical Student”*

Ask a question, get a thoughtful answer — but only when asked.

---

*ChatGPT for patient education, LLM literature synthesis, draft notes*



## Agentic AI

*“The Senior Resident”*

You give it a goal — it plans, uses tools, returns a draft for your sign-off.

---

*Living guidelines, agentic trial matching, autonomous evidence synthesis*



*Each tier is already in practice today. The frontier is moving — quickly — toward the right.*

# From Clinic to Computation

*The same clinical problems, named as the ML problems you already know*



## Treatment Selection

*Policy learning / ITE*

Pick the action that maximizes this patient's outcome.

---

*COMPASS-CSPC, causal forests, T-/X-learners*



## Staging & Imaging

*Detect / segment / quantify*

Find it, outline it, measure it on MRI / PET / path.

---

*PSMA-PET quantification, MRI lesion detection, Gleason grading*



## Evidence at Scale

*Retrieval + NLP + meta-analysis*

Keep a synthesis current as new trials report.

---

*Living guidelines, LLM screening & extraction*



## Genomic Risk

*Representation learning on omics*

Turn high-dimensional omics into a usable signal.

---

*Decipher, ctDNA classifiers, omic foundation models*



*Every clinical task in this talk maps to a problem class you can model, validate, and improve.*

# AI Voice Capabilities

*OpenAI — Natural Voice Interaction*



Natural, real-time voice conversation with AI. The foundation for ambient documentation, patient communication, and hands-free clinical decision support.

Video: <https://www.youtube.com/watch?v=c2DFg53Zhvw>

# AI Vision Capabilities

*OpenAI — Multimodal Visual Understanding*



AI interprets real-world scenes, reads medical images, and understands visual context in real time. Imagine this applied to radiology, pathology slides, or bedside assessment.

Video: <https://www.youtube.com/watch?v=RI-BxtCx32s>

# Ambient Documentation in Practice

*Microsoft DAX Copilot — AI Scribe Demo*



DAX Copilot listens to the patient encounter, generates the clinical note in real time, and drafts orders — already deployed at Mayo Clinic and 600+ health systems.

# Three Axes of AI in Clinic

*Same underlying shift — three different surfaces where it shows up*



**AXIS 1**

## Patient Journey

*AI in the disease.*

From pre-cancer risk to end-of-life — touching diagnosis, staging, treatment selection, and monitoring.

---

Artera · COMPASS · PSMA-AI · ctDNA



**AXIS 2**

## Clinical Workflows

*AI in the system.*

Removing friction from the work of cancer care — intake, documentation, decision support, and trial operations.

---

Ambient scribes · MAYA · Living guidelines



**AXIS 3**

## Patient–Clinician Interaction

*AI in the relationship.*

Patients are already using AI to engage with their cancer — interpretation, education, trial-finding, communication.

---

Open Notes · Trial matching · Virtual counselors



We'll walk through each axis in turn — anchored on one prostate cancer patient: Mr. K.



# AI in Clinical Workflows

*Where AI quietly removes friction from the work of cancer care*

## 1 Intake & Triage Automation

AI intake assessments, automated routing of new referrals to the right specialist.

## 2 AI-Assisted Chart Search

Tools like MAYA surface relevant prior notes, imaging, labs in seconds — not hours of scrolling.

## 3 Ambient Documentation

DAX, Abridge, Carellon — listen, draft note, propose orders. Returning time to the encounter.

## 4 Decision Support

Guideline-concordant recommendations surfaced at point of care — and (next slide) updated continuously.

## 5 Trial Operations

Activation, patient identification, eligibility screening, reporting — automation across the trial lifecycle.

## 6 Registry & Database Automation

Cancer registries, real-world data extraction, structured reporting — without armies of abstractors.



*Most workflow AI is invisible to patients. The deep dive: living evidence — the only one they will eventually feel.*

# AI Across the Prostate Cancer Journey

AXIS 1 | JOURNEY

*From pre-cancer to end-of-life — AI is already at every stop*

Pre-cancer	Diagnosis	Staging	Treatment	Follow-up	End-of-Life
Population PSA risk modeling	MRI lesion detection	PSMA PET interpretation	Treatment selection (Artera)	PSA trajectory modeling	Symptom prediction
Genetic risk (HOXB13, BRCA)	Pathology foundation models	Automated quantification	Individualized policy (COMPASS)	ctDNA monitoring	Care preference prediction
AI-guided MRI screening	Augmented Gleason grading	Multi-modal prognostication	Dose personalization in RT	Recurrence prediction	Advance care planning
Lifestyle / EHR risk	Multi-modal diagnostics	LN prediction (vs ePLND)	Toxicity prediction	Survivorship support	Caregiver support



*Two deep dives ahead — Artera at the moment of treatment selection, COMPASS for individualized treatment policy.*



# Meet Mr. K

## 62-year-old man — newly diagnosed prostate cancer

PSA 12 ng/mL | Gleason 4+3 = 7 (Grade Group 3) | cT2c

MRI: PI-RADS 5 lesion, no extracapsular extension

Conventional staging: no metastatic disease

*Unfavorable intermediate risk. The decision in front of us:*

## Definitive radiation alone — or radiation plus short-course ADT?

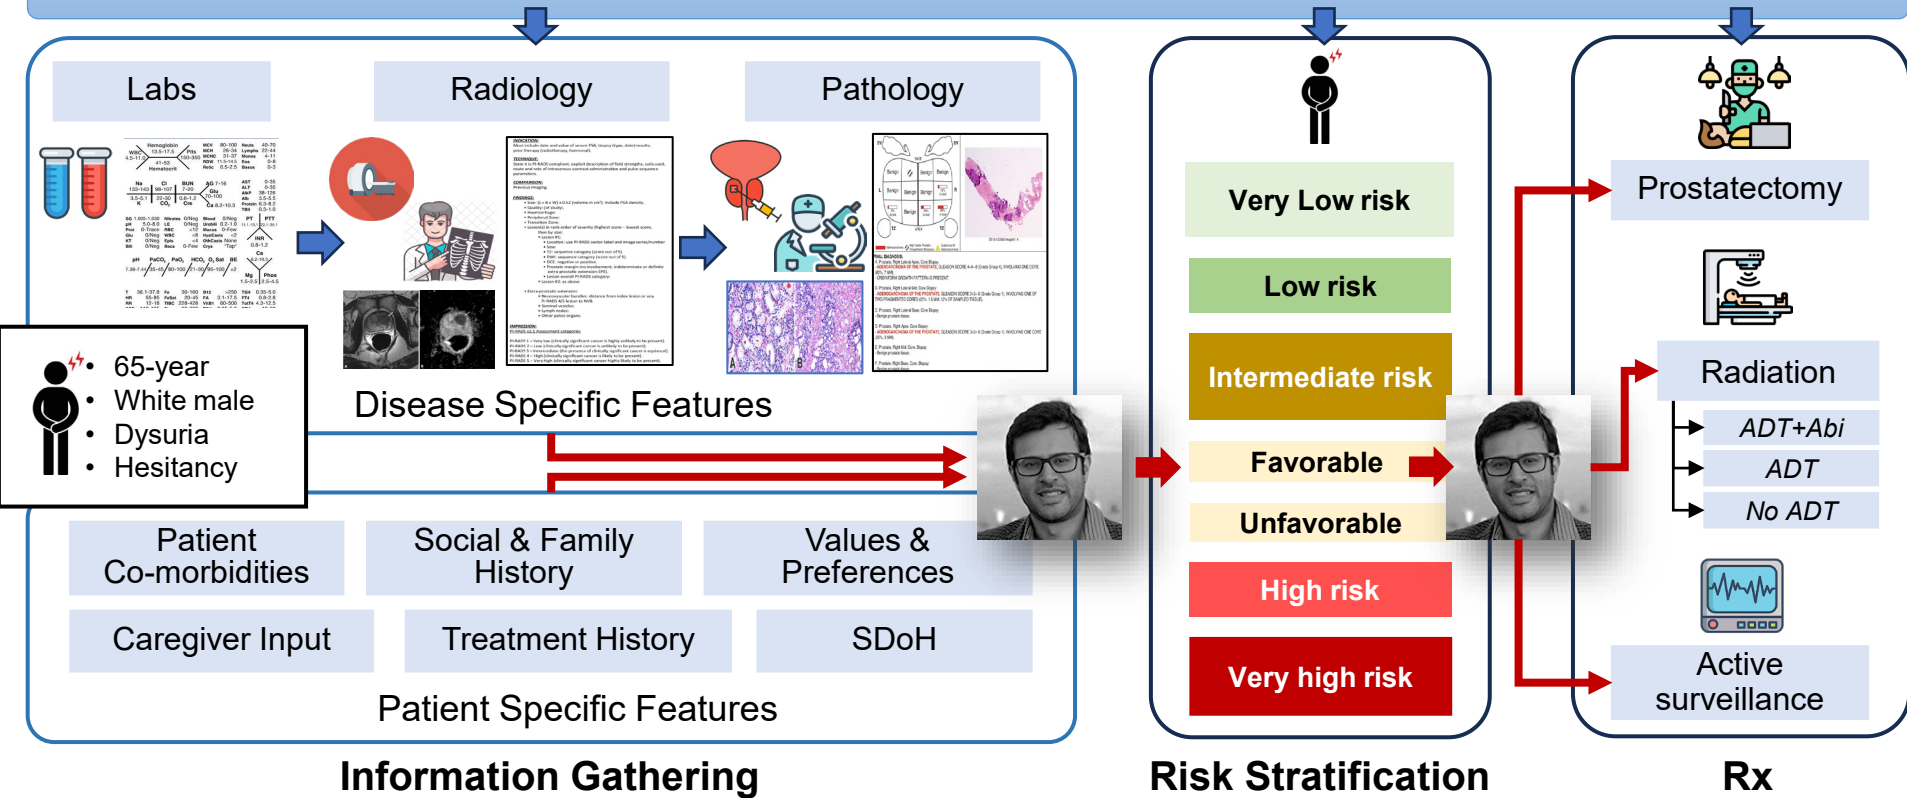
Trial-average benefit of ADT in this group: small.

Individual benefit: highly variable. How do we choose?



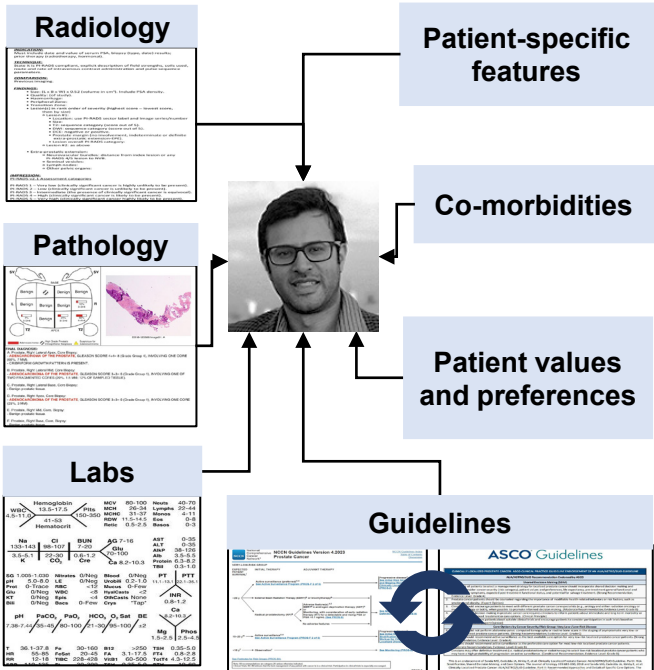
# Process of treatment selection: localized prostate cancer

## Prostate Cancer Guidelines



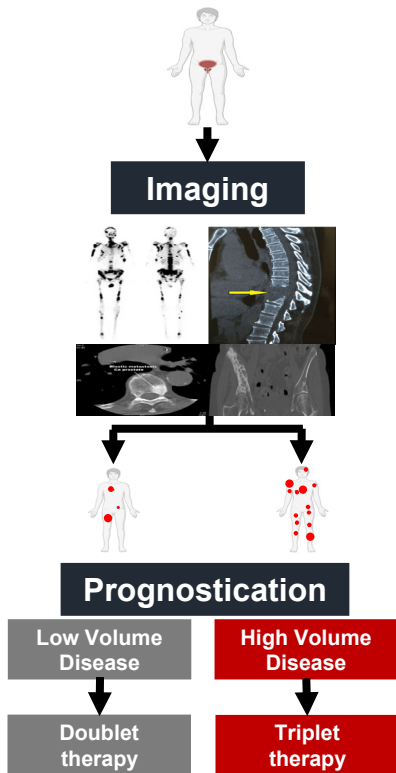
# Challenges in Treatment Selection and Predicting Outcomes

## 1. Too much to do !!

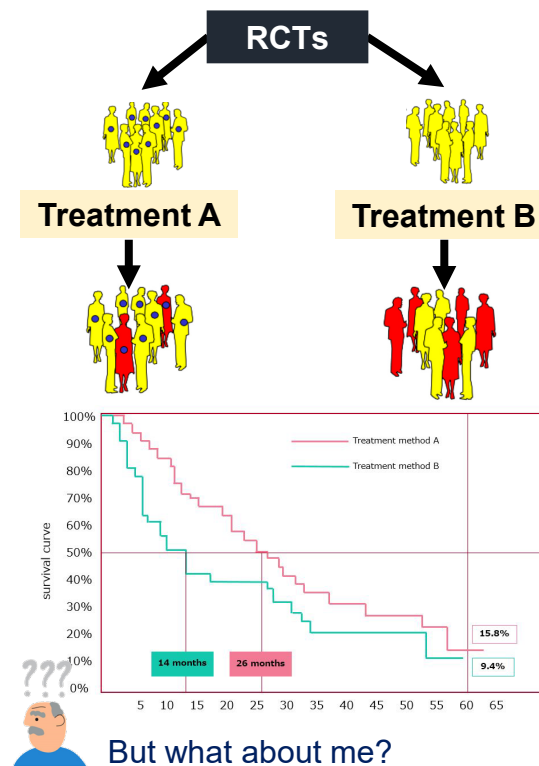


## 2. Rapidly changing evidence

## 3. No real biomarkers



## 4. RCTs report average results

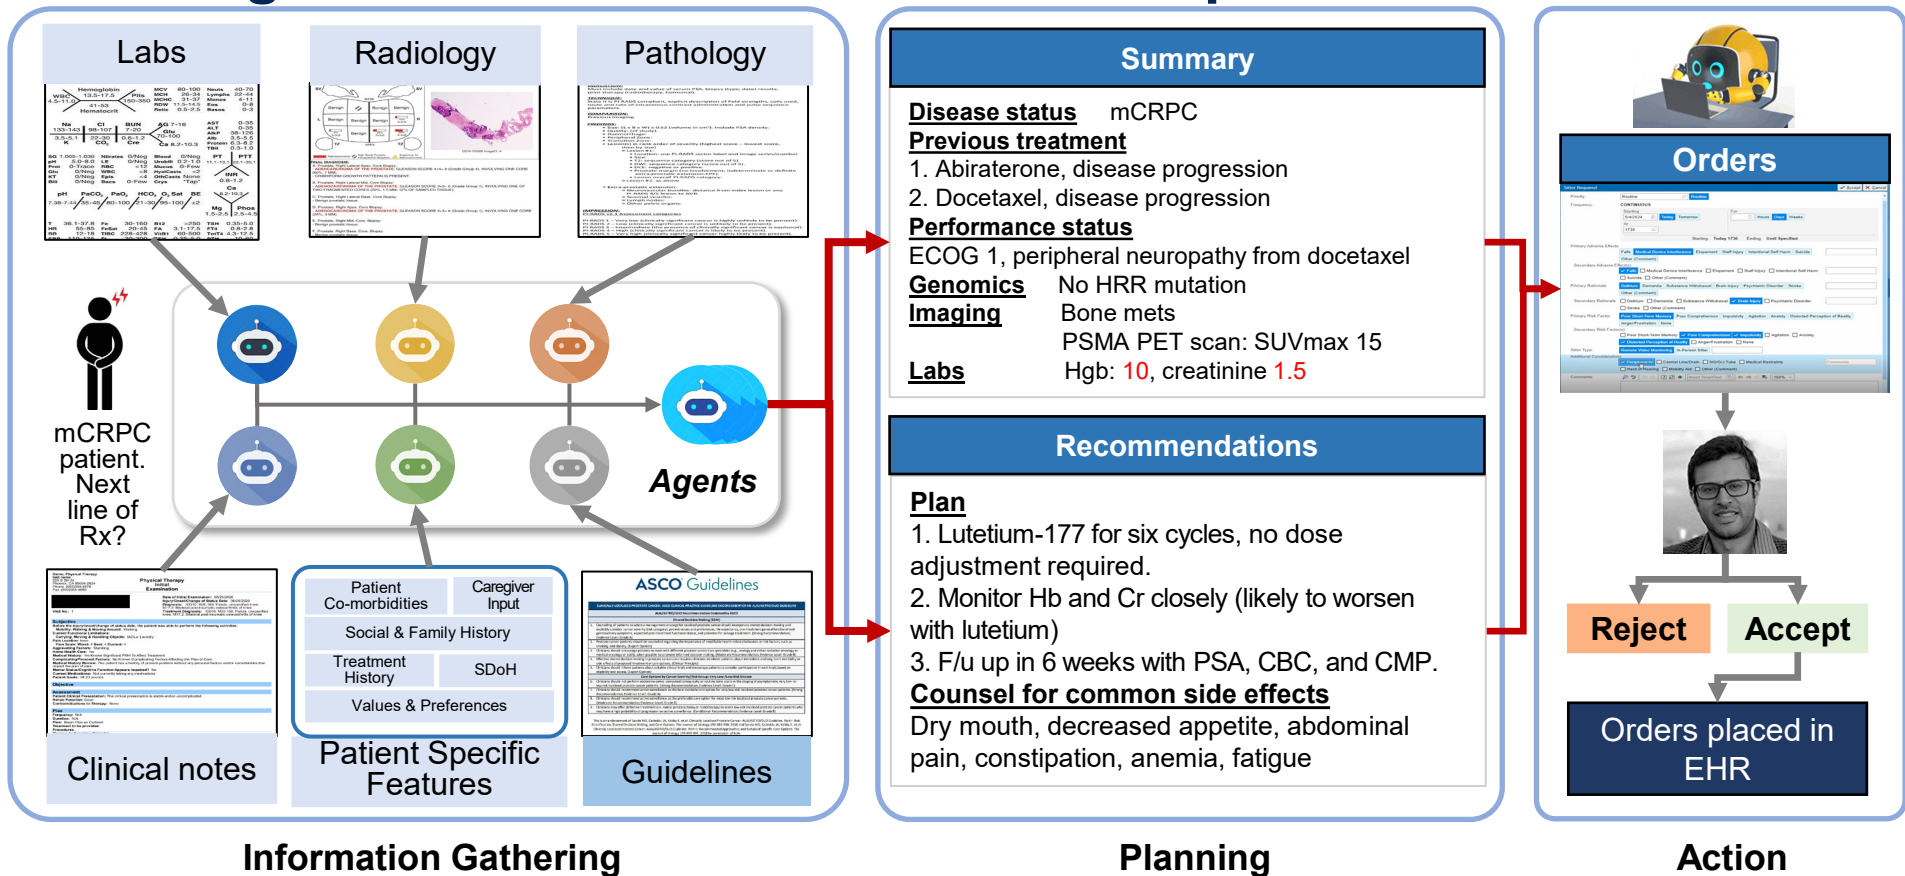




# The Modern AI Clinic

- AI agents acting as assistants to optimize treatment decisions
- Artificial intelligence-enabled **living clinical practice guidelines**
- **Patient-specific (individualized) treatment effects**
- **Digital biomarkers** utilizing multi-modal data

# #1 AI agents as assistants to healthcare providers





# The Modern AI Clinic

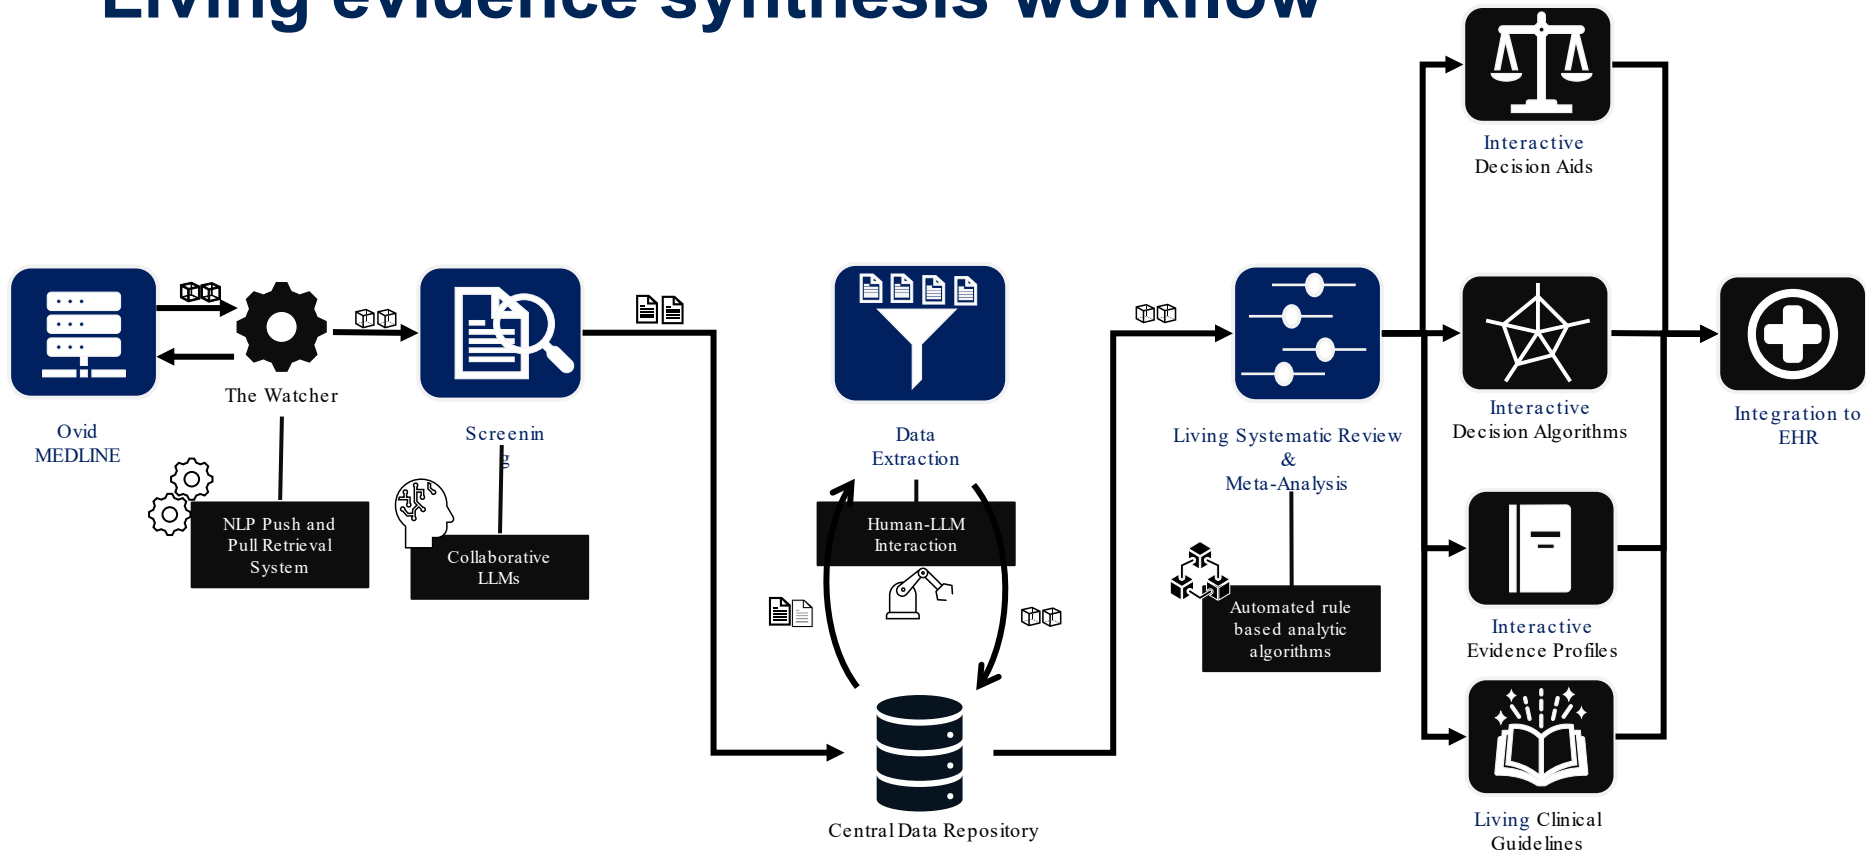
- **AI agents** acting as assistants to optimize treatment decisions
- **Artificial intelligence-enabled living clinical practice guidelines**
- **Patient-specific (individualized) treatment effects**
- **Digital biomarkers** utilizing multi-modal data







## #2 AI-enabled Living clinical practice guidelines

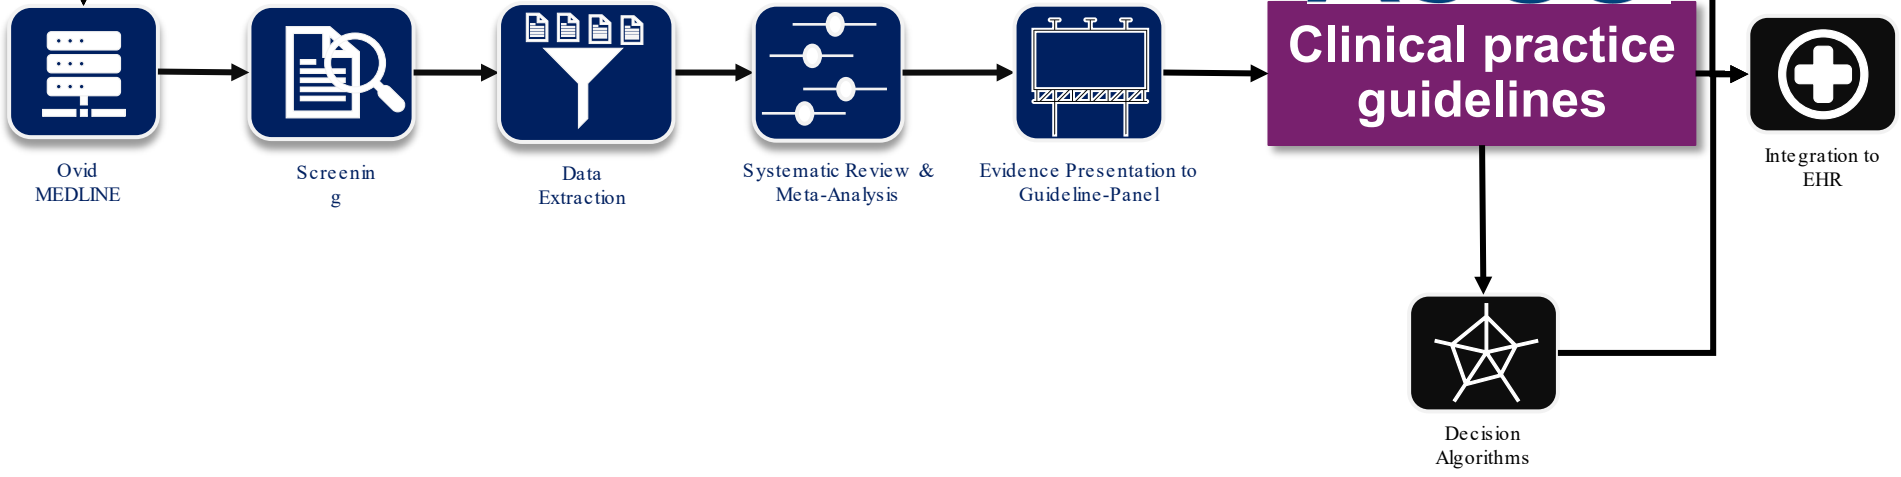
“Optimization of the standard guideline process, such that it allows updating of individual recommendations as soon as new relevant evidence becomes available”

# Living evidence synthesis workflow





-  GETUG  
NCT00104715
-  CHARTED  
NCT00309985
-  STAMPEDE  
NCT00268476
-  LATITUDE  
NCT01715285



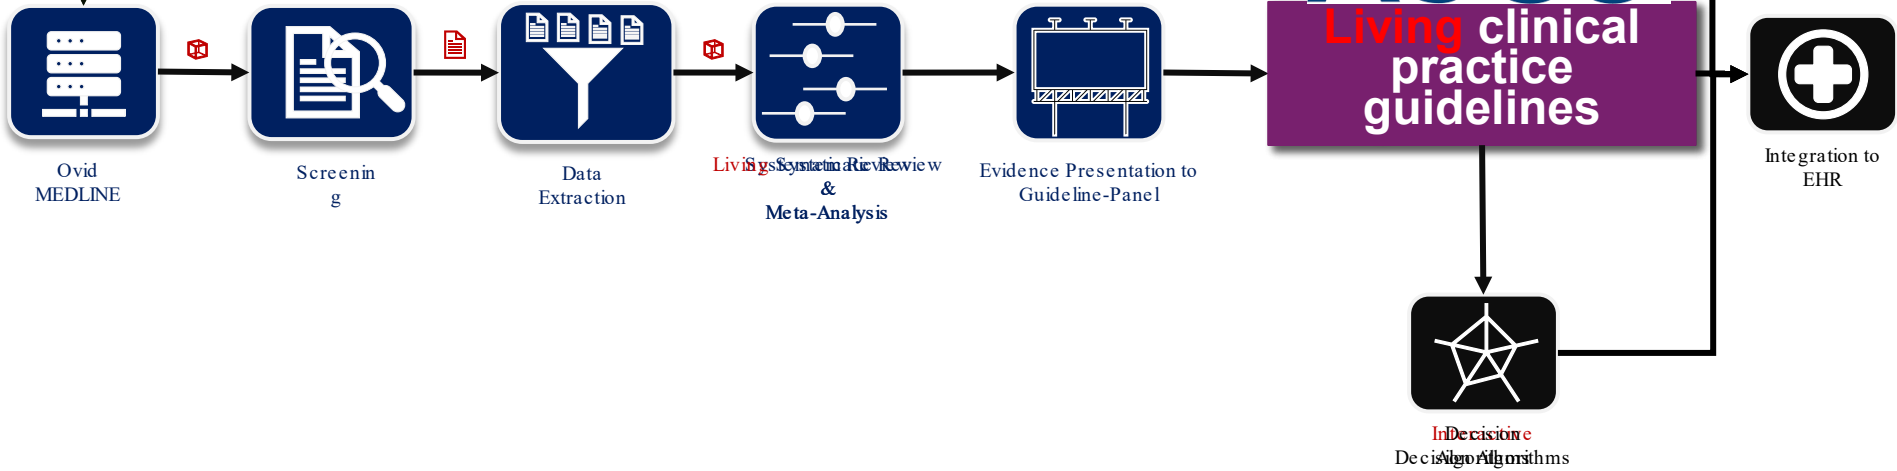


 GETUG  
 NCT00104715  
 CHAARTED  
 NCT00309985  
 STAMPEDE  
 NCT00268476  
 LATITUDE  
 NCT01715285

 TITAN  
 NCT02489318

 ENZAMET  
 NCT02446405

 ARCHES  
 NCT0267896



# Living interactive evidence synthesis framework has transformed the evidence ecosystem for prostate cancer

## Real-time evidence synthesis

10+ abstracts & full-text publication

JAMA Oncology | Original Investigation

### First-line Systemic Treatment Options for Metastatic Castration-Sensitive Prostate Cancer A Living Systematic Review and Network Meta-analysis

Irbaz Bin Riaz, MD, PhD; Syed Arsalan Ahmed Naqvi, MBBS; Huan He, PhD; Noureen Asghar, MBBS;

EUROPEAN UROLOGY 87 (2025) 626–640

available at [www.sciencedirect.com](http://www.sciencedirect.com)  
journal homepage: [www.europeurology.com](http://www.europeurology.com)



Review – Prostate Cancer  
Elena Castro, David Lorente, David Olmos on pp. 641–642 of this issue

### Heterogeneity of the Treatment Effect with PARP Inhibitors in Metastatic Castration-resistant Prostate Cancer: A Living Interactive Systematic Review and Meta-analysis

Syed Arsalan Ahmed Naqvi<sup>a,c,\*</sup>, Irbaz Bin Riaz<sup>a,c,\*</sup>, Arifa Bibi<sup>b</sup>, Muhammad Ali Khan<sup>a</sup>,

### Systemic treatment options for metastatic castration resistant prostate cancer: A living systematic review

Syed Arsalan Ahmed Naqvi<sup>a,c,\*</sup>, Muhammad Umair Anjum<sup>a,c,\*</sup>, Arifa Bibi<sup>b</sup>, Muhammad Ali Khan<sup>a</sup>, Kaneez Zahra Rubab Khakwani<sup>c</sup>, Huan He<sup>d</sup>, Manal Imran<sup>c</sup>, Syeda Zainab Kazmi<sup>c</sup>, Ammad Raina<sup>f</sup>, Ewan K Cobran<sup>e</sup>, R Bryan Rumble<sup>h</sup>, Thomas K Oliver<sup>h</sup>, Neeraj Agarwal<sup>i</sup>, Yousef Zakharia<sup>a</sup>, Mary-Ellen Taplin<sup>j</sup>, Oliver Sartor<sup>k</sup>, Parminder Singh<sup>a</sup>, Jacob J Orma<sup>k</sup>, Daniel S Childs<sup>k</sup>, Rahul A Parikh<sup>l</sup>, Rohan Garje<sup>e</sup>, Mohammad Hassan Murad<sup>n</sup>, Alan H Bryce<sup>o</sup>, Irbaz Bin Riaz<sup>a</sup>

## Methods innovation

Journal of the American Medical Informatics Association, 2025, 30(4), 638–647  
<https://doi.org/10.1093/jamia/ocae329>  
Advance access publication 21 January 2025  
Research and Applications



Research and Applications

### Collaborative large language models for automated data extraction in living systematic reviews

Muhammad Ali Khan, MBBS<sup>1</sup>, Umair Ayub, PhD<sup>1</sup>, Syed Arsalan Ahmed Naqvi, MD<sup>1</sup>, Kaneez Zahra Rubab Khakwani, MD<sup>2</sup>, Zaryab bin Riaz Sipra, MD<sup>3</sup>, Ammad Raina, DO<sup>4</sup>, Sihan Zhou, BS<sup>1</sup>, Huan He<sup>o</sup>, PhD<sup>5</sup>, Amir Saedi, MS<sup>1,6</sup>, Bashar Hasan, MD<sup>7</sup>, Robert Bryan Rumble, MSc<sup>8</sup>, Danielle S. Bitterman, MD<sup>9</sup>, Jeremy L. Warner, MD, MS<sup>10,11,12,13</sup>, Jia Zou, PhD<sup>4</sup>, Amye J. Tevaarwerk, MD<sup>14</sup>, Konstantinos Leventakos, MD, PhD<sup>7</sup>, Kenneth L. Keil, MD, MPH<sup>15</sup>, Jeanne M. Palmer, MD<sup>1</sup>, Mohammad Hassan Murad, MD<sup>7</sup>, Chitta Baral, PhD<sup>6</sup>, Irbaz bin Riaz, MD, PhD<sup>1,16</sup>



Journal of Biomedical Informatics

Available online 28 May 2025, 104860

In Press, Journal Pre-proof What's this?

### Do it faster with PICOS: Generative AI-Assisted systematic review screening

Sai Krishna Vallamchethla<sup>a</sup>, Omar Abdelkader<sup>a</sup>, Ali Elnaggar<sup>b</sup>, Doaa Ramadan<sup>a</sup>, Md Manjurul Islam Shourav<sup>c</sup>, Irbaz B. Riaz<sup>c</sup>, Michelle P. Lin<sup>a</sup>,



MAYO CLINIC PROCEEDINGS:  
DIGITAL HEALTH

### Future of Evidence Synthesis: Automated, Living, and Interactive Systematic Reviews and Meta-analyses

Irbaz Bin Riaz, MD, MS, MBI, PhD; Syed Arsalan Ahmed Naqvi, MD; Bashar Hasan, MD; and Mohammad Hassan Murad, MD, MPH

## Living guidelines

ASCO Special Articles



### Systemic Therapy in Patients With Metastatic Castration-Resistant Prostate Cancer: ASCO Guideline Update

Rohan Garje, MD<sup>a</sup>, Irbaz Bin Riaz, MD, PhD<sup>b</sup>, Syed Arsalan Ahmed Naqvi, MD<sup>c</sup>, R. Bryan Rumble, MSc<sup>d</sup>, Mary Ellen Taplin, MD, FASCO<sup>e</sup>, Terry M. Kung'u, MBA<sup>f</sup>, Daniel Herchenhrom, MD, PhD<sup>g</sup>, Tian Zhang, MD, MI<sup>h</sup>, Kathryn E. Beckermann, MD, PhD<sup>i</sup>, Neha Vapiwala, MD<sup>j</sup>, Michael A. Carducci, MD, FASCO<sup>k</sup>, Paul Celisno, MD, FASCO<sup>l</sup>, Sebastian J. Horst, MD, MSc<sup>m</sup>, Anubh Bassi, MBBS, MPH<sup>n</sup>, Hala Bono, MSc<sup>o</sup>, Alan H. Bryce, MD<sup>p</sup>, Peng Wang, MD, PhD<sup>q</sup>, Elizabeth Wulf-Burdick, MD<sup>r</sup>, Lisa Bordes, PhD, MD<sup>s</sup>, Andrew Loblaw, MD, FASCO<sup>t</sup>, Robert J. Hamilton, MD, MPH<sup>u</sup>, Hamid Emamekhoh, MD<sup>v</sup>, Thomas A. Hope, MD<sup>w</sup>, Huan He, PhD<sup>x</sup>, M. Hassan Murad, MD, MPH<sup>y</sup>, Hongfang Liu, PhD<sup>z</sup>, James Elbert Williams Jr, MSc<sup>aa</sup>, and Rahul A. Parikh, MBBS, PhD<sup>ab</sup>

DOI: <https://doi.org/10.1200/JCO.2025.0007>

# ASCO

AMERICAN SOCIETY OF CLINICAL ONCOLOGY

## ASCO Genitourinary Cancers Symposium

# ESMO

GOOD SCIENCE  
BETTER MEDICINE  
BEST PRACTICE

# AMIA

INFORMATICS PROFESSIONALS. LEADING THE WAY.



# The Modern AI Clinic

- **AI agents** acting as assistants to optimize treatment decisions
- Artificial intelligence-enabled **living clinical practice guidelines**
- **Patient-specific (individualized) treatment effects**
- **Digital biomarkers** utilizing multi-modal data



# Treatment: A Decision Made on Average

*The mHSPC question:*

**ADT alone, doublet (ADT + ARPI), or triplet (ADT + ARPI + docetaxel)?**



## Trial-Average Effects

PEACE-1, ARASENS, ENZAMET, TITAN, STAMPEDE — all show benefit on average.  
None tells us which individual patient benefits.



## Underpowered Subgroups

High-volume vs. low-volume, de novo vs. metachronous — subgroup forests are noisy.  
We act on them anyway.



## Where Experts Disagree

The live controversies are exactly the patients the trial average can't resolve — where an individualized estimate would change the call.

# COMPASS-CSPC

*ML for individualized treatment policy in mHSPC*

## THE APPROACH

1

### Longitudinal EHR cohort

Multi-site mCSPC patients with full treatment + outcome trajectories

2

### ML treatment policy learning

Estimate individualized treatment effects across doublet / triplet / ARPI options

3

### Clinician-facing dashboard

Decision support — not replacement. Surface the model's reasoning; clinician decides.

## THE BIGGER FRAME

*“The next mHSPC debate won't be settled by trial subgroups. It will be settled by individualized policy models.”*



*Mr. K in 18 months: if his disease progresses, the question stops being 'what's the trial-average best.' It becomes 'what's best for him.'*



# COMPASS-CSPC Under the Hood

*Estimating individualized treatment effects from observational EHR data*

1

## Potential Outcomes

We want each patient's outcome under every option — doublet, triplet, ARPI alone. Only one is ever observed; the rest are counterfactual. The estimand is the contrast between them — the individualized treatment effect.

2

## Why EHR Is Not an RCT

Treatment was never randomized. Sicker patients get intensified therapy (confounding by indication); some patient types never receive a given option (positivity). A naive outcome comparison recovers bias, not effect — adjustment is the whole game.

3

## Validating the Counterfactual

You can't directly score a prediction you never observe. We lean on held-out sites, calibration against known trial subgroup effects, and sensitivity analysis for unmeasured confounding — not accuracy alone. Decision support stays human-interrogable.


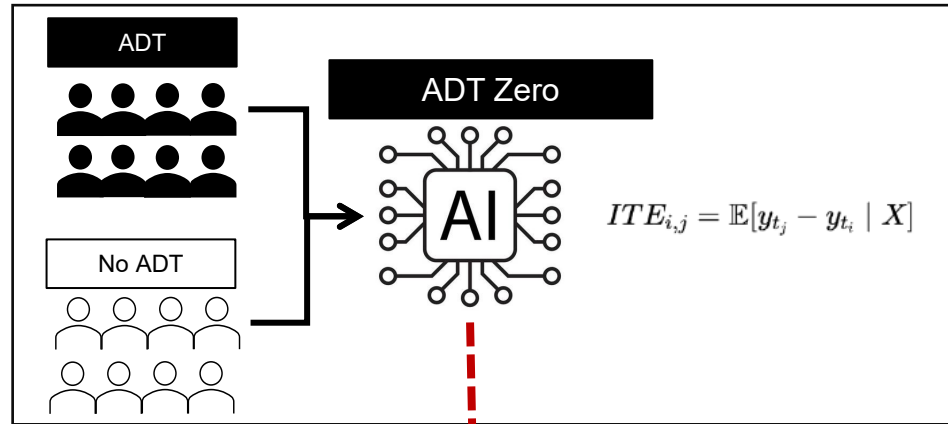


# The Modern AI Clinic

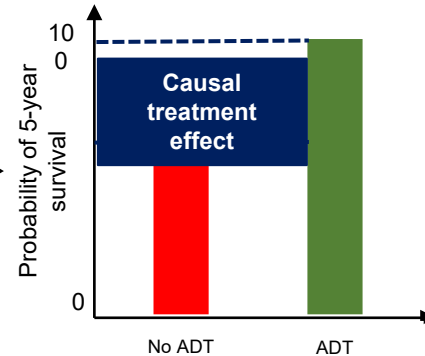
- **AI agents** acting as assistants to optimize treatment decisions
- Artificial intelligence-enabled **living clinical practice guidelines**
- **Patient-specific (individualized) treatment effects**
- **Digital biomarkers** utilizing multi-modal data

# #3 Individualized treatment effects

## Causal inference artificial intelligence



Prostate cancer patient



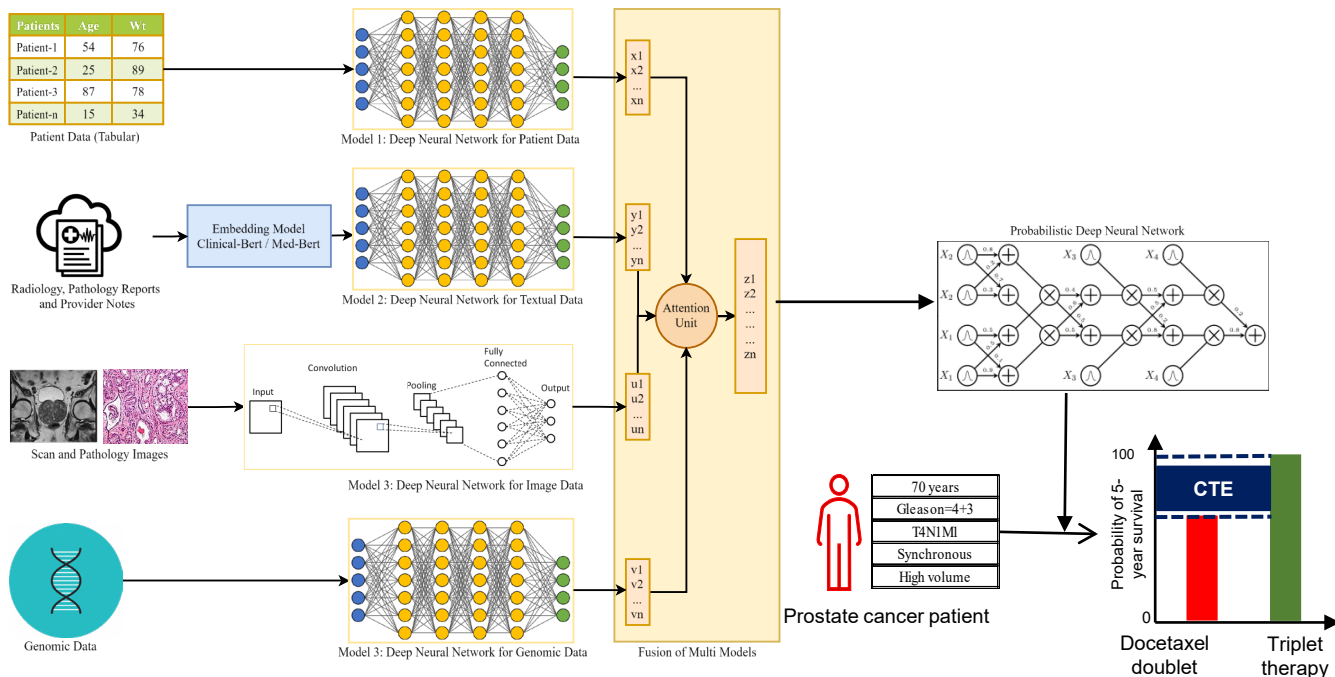
# Multimodal individualized treatment effects

## Probabilistic modeling

- To develop a multi-modal artificial intelligence model that integrates longitudinal data:

- Clinical
- Pathological
- Imaging
- Genomic data

- Provides patient-specific treatment effects using a probabilistic deep-learning causal inference framework for prostate cancer patients



# AI in the Patient–Clinician Interaction

AXIS 3 | INTERACTION

*Patients are already using AI to engage with their cancer care — clinicians need a point of view*

## 1 Know the Doctor

Patient-facing physician matching — finding the right oncologist for the right patient.

## 2 Consult Prep

Pre-visit preparation — particularly relevant for high-stakes prostate cancer treatment decisions.

## 3 Virtual Treatment Companion

Patient-facing education on complex regimens — ARSI, chemo, radioligand therapy.

## 4 Virtual Genetic Counselor

BRCA, HRR, Lynch — bridging the gap between testing demand and counselor capacity.

## 5 AI-Assisted Inbox & Open Notes

Patients reading their PSA, MRI, and pathology reports — with LLMs interpreting in plain language.

## 6 Patient-Facing Trial Matching

Community outreach for trial access — from health-system bottleneck to patient-initiated discovery.



*The deep dive ahead: trial matching is the bridge — clinician-facing in clinic, patient-facing at home.*



# Clinical Trial Matching (CTM)

*Bridging the access gap — from 'what trials exist' to 'which trial fits this patient'*

## THE GAP

< 8%

*of adult cancer patients enroll in a clinical trial*

The bottleneck isn't trial availability — it's the work of matching the right patient to the right trial.

## BEACON CTM

### Patient facing trial matching



#### Tiered Matching

Ranks trials by clinical fit, not just keyword overlap



#### Eligibility Explainability

Surfaces why a patient is or isn't eligible — at the criterion level



#### Patient Profile Panel

Clinician sees the structured profile the model is reasoning over



#### Dual-Mode Chatbot

Quick lookup or guided clinician/patient conversation



*If Mr. K progresses to mCRPC, CTM is the difference between 'no trial available' and 'three trials, ranked, with eligibility flagged.'*



# What to Watch Out For

*Three tensions the field has not yet resolved*

1

## The Evidence-Generation Gap

AI tools are being deployed faster than they're being validated. The field has no agreed standard for what counts as sufficient evidence — RCT? Prospective registry? External validation across N sites? Until this is resolved, every AI tool lives in a credibility limbo.

2

## Distribution Shift & External Validation

Almost every model here was trained on US/European data. “Does this work on my patients?” is empirical, not rhetorical — demographics, biology, treatment patterns, and label definitions all shift across sites. Most papers report internal performance; few test external validation under distribution shift.

3

## Division of Labor — What Should Be Automated?

The useful question isn't replacement — it's which tasks to automate and which to keep human. Pattern recognition, documentation, surveillance, and evidence retrieval are increasingly learnable. Judgment under uncertainty, accountability, and the goals-of-care conversation are not — and may not be the right targets. If you build these systems, you draw that line: design the human-AI handoff deliberately, not by default.

# Back to Mr. K

*What AI touches in his care today, in 2 years, and what it shouldn't touch at all*



## TODAY

*What AI already touches*

- Pathology foundation models confirm Gleason grade
- Artera predicts his benefit from adding ADT
- AI-quantified MRI informs RT planning



## IN 24 MONTHS

*What AI will likely touch*

- COMPASS-style policy models guide intensification at progression
- PSMA PET interpretation fully AI-augmented
- Living guidelines surface relevant trials in real time



## NEVER

*What AI shouldn't touch*

- The conversation about goals of care
- What this diagnosis means for him and his family
- The decision itself — it remains his and ours



*AI in uro-oncology is neither as far along as the hype, nor as distant as the skeptics claim. The interesting work is in the middle.*



# Thank You

Questions & Discussion

---

**Irbaz Bin Riaz, MBBS, MS, MBI, PhD**

Director, Cancer AI | Director, BEACON Lab

Mayo Clinic Arizona

[riaz.irbaz@mayo.edu](mailto:riaz.irbaz@mayo.edu) | [@IrbazRiaz](https://twitter.com/IrbazRiaz) | [beaconlab.ai](https://beaconlab.ai)