Computational
PRECISION HEALTH

UC Berkeley Center for Healthcare
Marketplace Innovation

Berkeley | UCSF

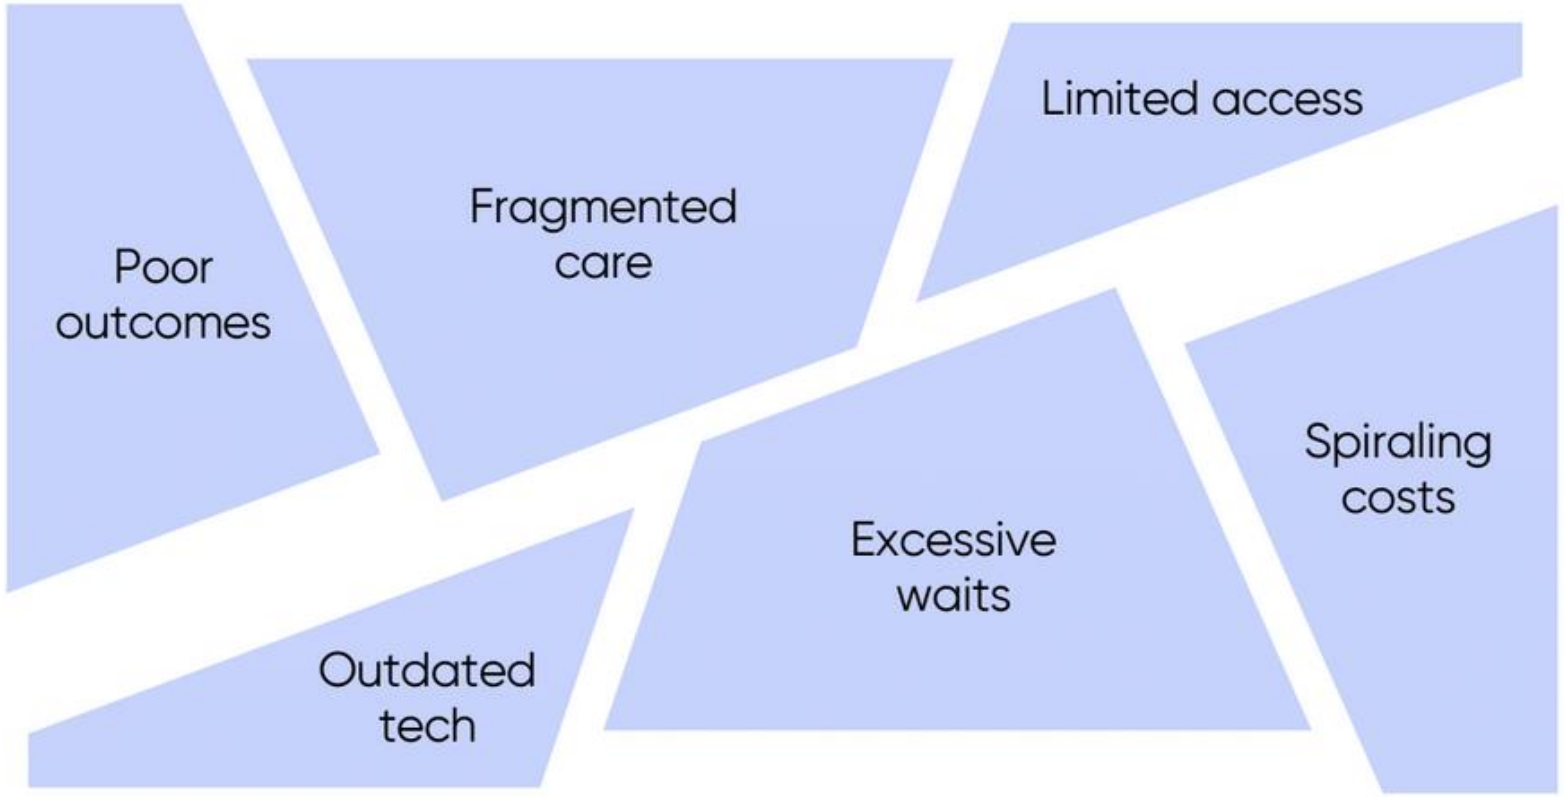
CTML
CENTER FOR TARGETED MACHINE
LEARNING AND CAUSAL INFERENCE

Causal AI for Health Care

Maya Petersen

CPH 100
November 13, 2026

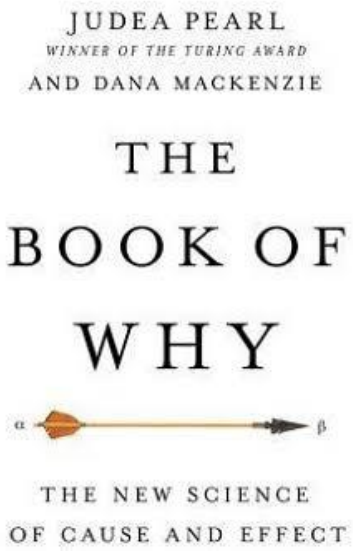
Today’s healthcare system is not sustainable

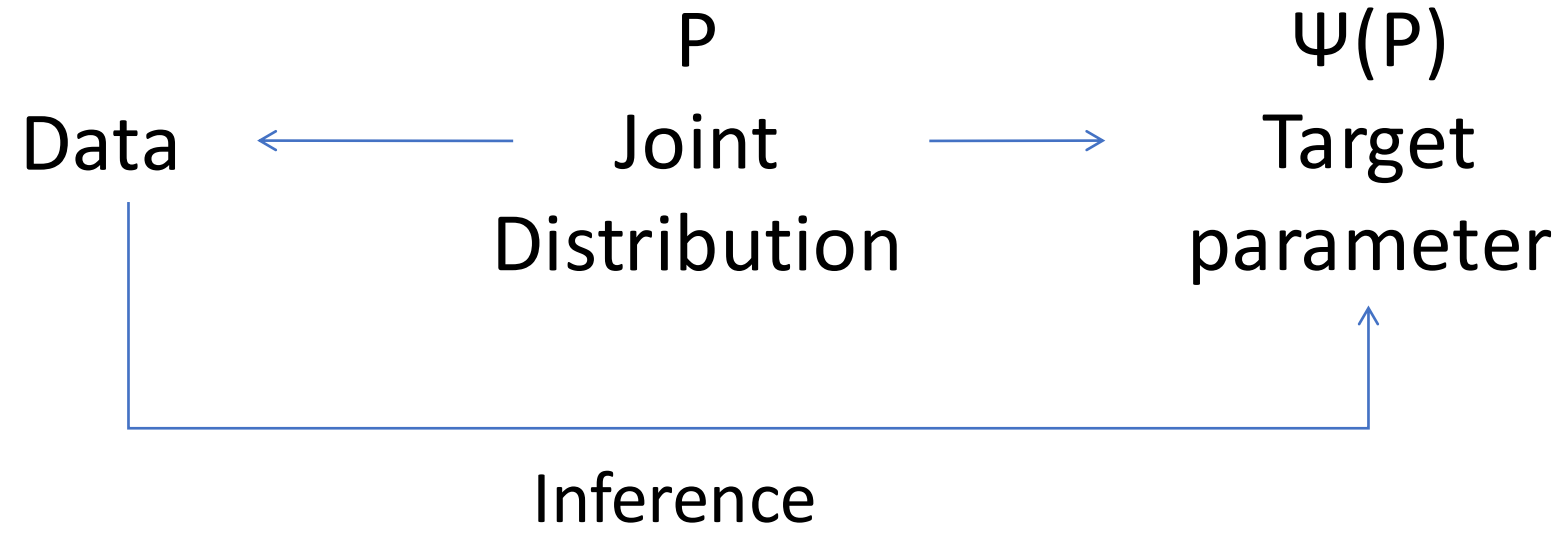


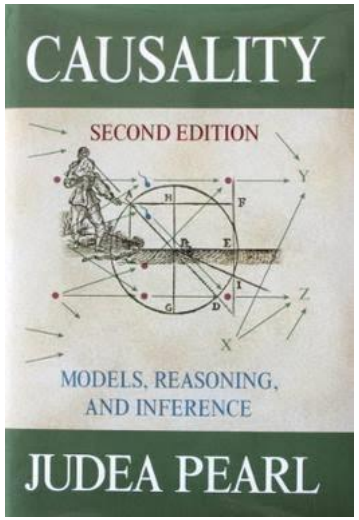
AI to the rescue?

- Causal inference required
- Need to rethink design and deployment of causal tools

Traditional Statistical Inference Paradigm



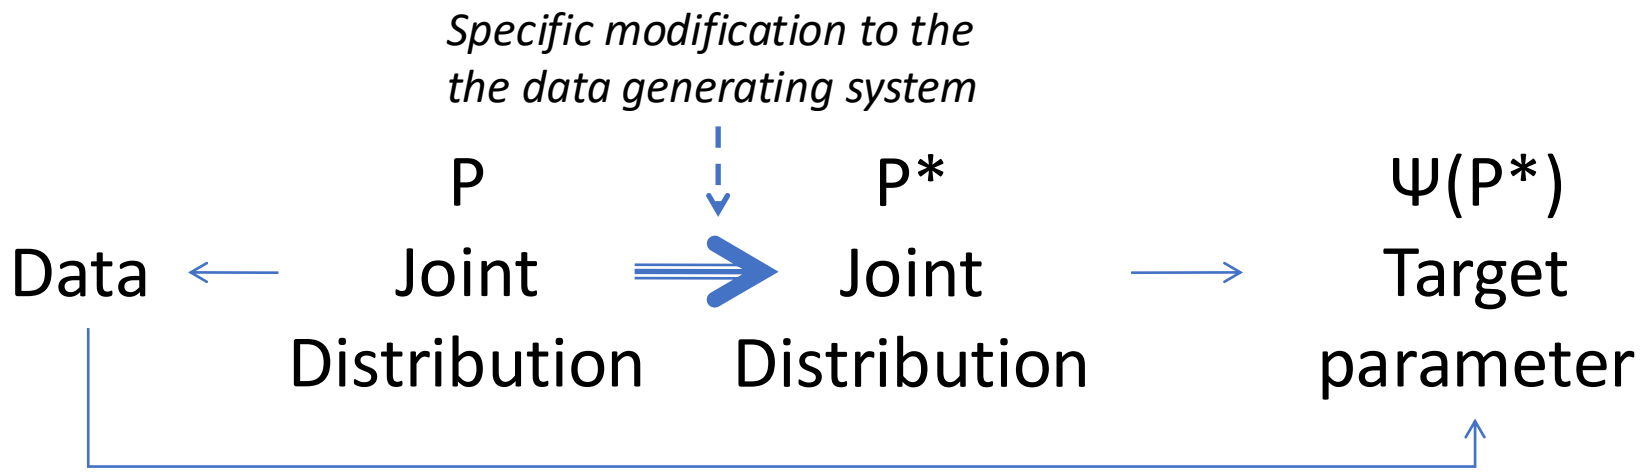




Example: What % of the population who had gastric bypass surgery went on to get diabetes?

Adapted from: cs.ucla.edu/~judea/jsm12

From a Statistical to Causal Analysis



Example: What % of the population would get diabetes if

- every eligible patient had surgery?
- none had surgery?
- only those with HA1c>6 had surgery?

Adapted from: cs.ucla.edu/~judea/jsm12

The causal question and the causal estimand

- Specify: the target population, ideal treatment strategies, time zero, outcome, and contrast
- Example:
  - Population & time zero: adults with BMI ≥ 35 and no diabetes, referral for surgery
  - Hypothetical Treatment strategies: gastric bypass within 6 months (t=1) versus no surgery (t=0); no loss to follow-up (c=0)
  - Outcome: type 2 diabetes within 3 years (Y)

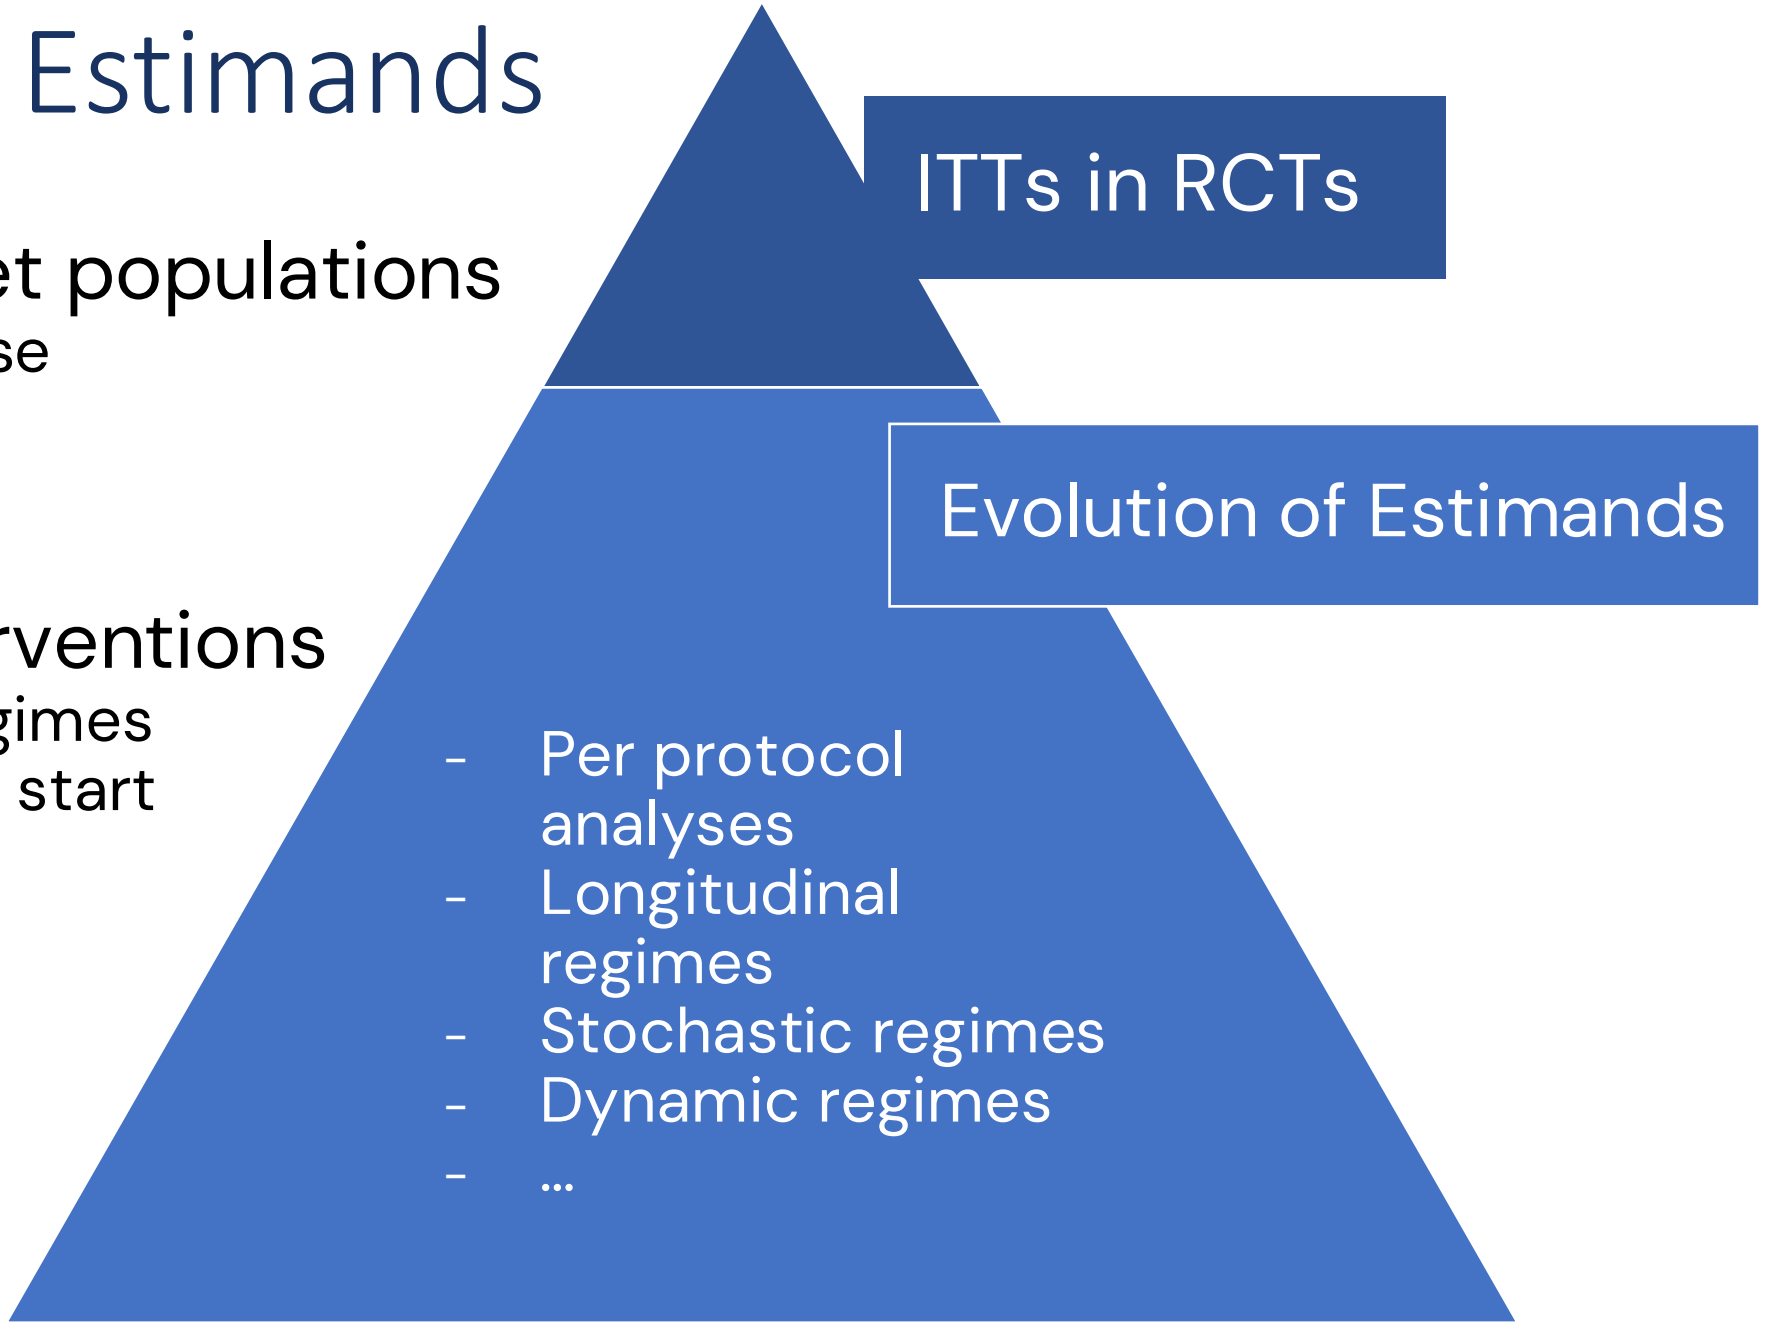
Potential outcomes: $\mathbf{Y_{t=1,\,c=0}}$ and $\mathbf{Y_{t=1,\,c=0}}$

Causal estimand: $\mathbf{\Psi^* = E(Y_{t=1,\,c=0}) - E(Y_{t=0,\,c=0})}$

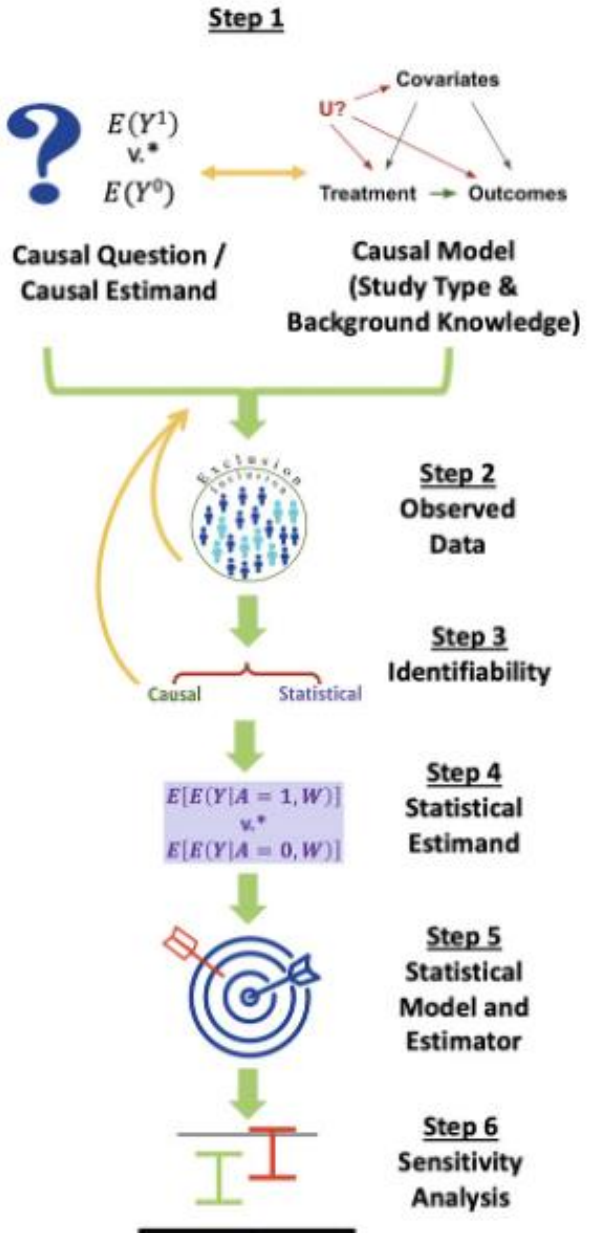
How would you modify this estimand if you wanted to understand how varying the A1C threshold for surgery would change risk of diabetes?

Classic Causal Estimands

- Restricted target populations
  - Ex. single disease
- Average effects
- One or few interventions
  - Ex. two drug regimes
  - Ex. threshold to start
- One or few:
  - Outcomes
  - Settings
  - Time points



The causal roadmap



- An explicit, itemized process for causal analyses
- Steps 1 – 4: Translate a causal question (framed as an ideal experiment) into a statistical estimation problem)
  - Language models
- Step 5: Statistical Estimation
  - Foundation/World models
- One framework across randomized trials, pragmatic trials, and fully observational analyses

Figure: Dang et al., Journal of Clinical and Translational Science, 2023

Agentic “Causal Co-Pilot”



Tianyue Zhou

Biostatistics



Yi Li

Biostatistics



Mark van der Laan

NEJM AI

Artificial Intelligence–Based Copilots to Generate Causal Evidence

Maya Petersen, M.D., Ph.D.,[1,2] Ahmed Alaa, Ph.D.,[2] Emre Kıcıman, Ph.D.,[3] Chris Holmes, Ph.D.,[4,5] and Mark van der Laan, Ph.D.[1]

CTML

CENTER FOR TARGETED MACHINE LEARNING AND CAUSAL INFERENCE

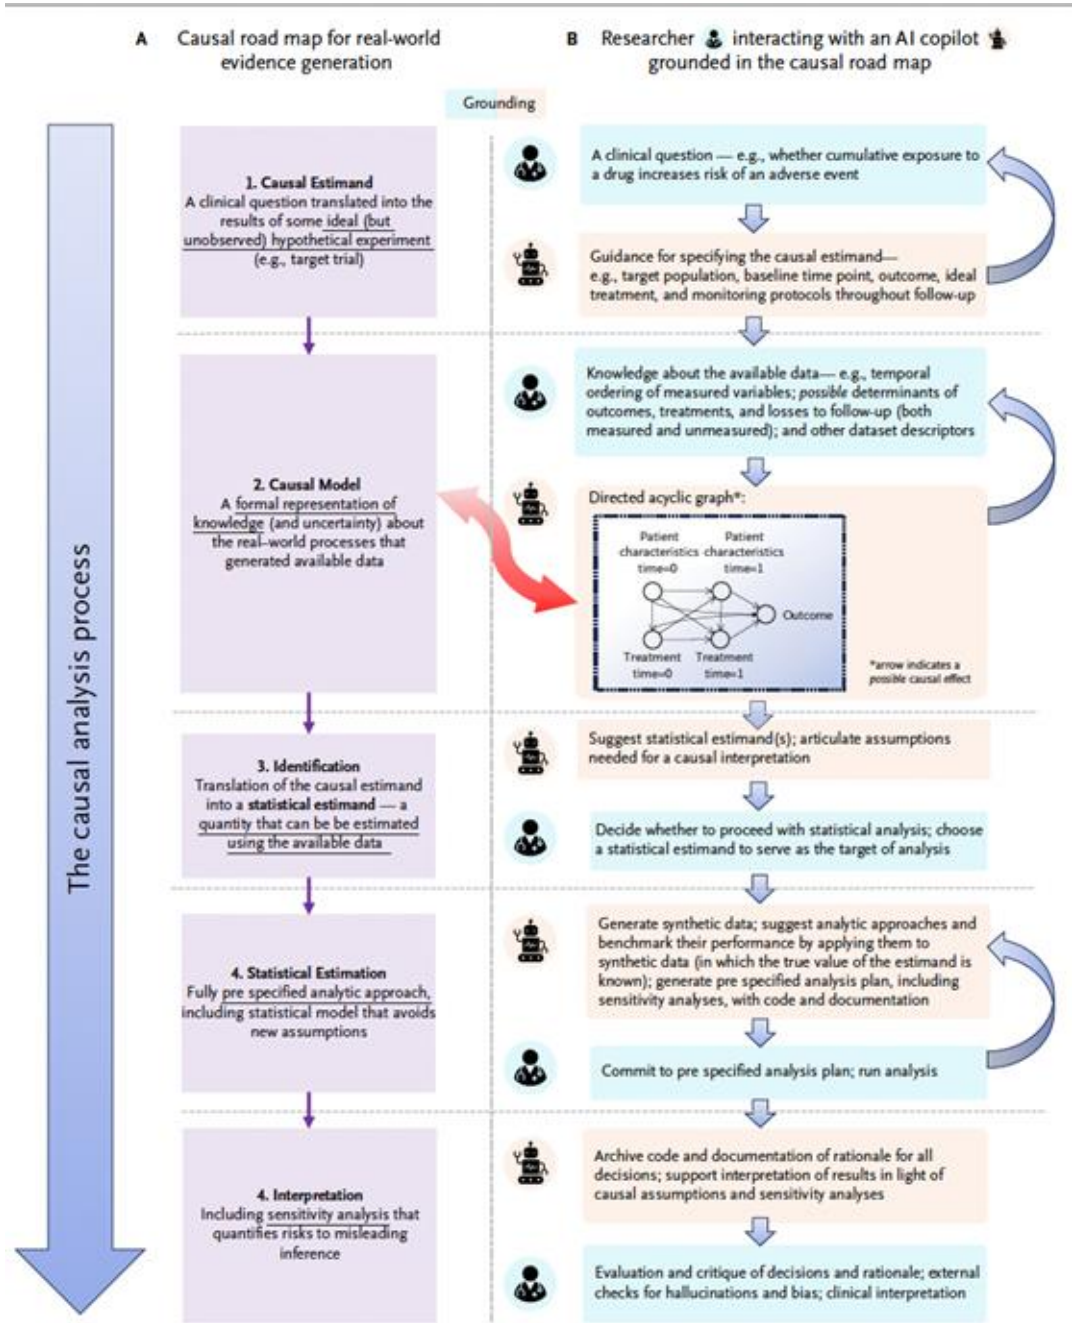
LLM-based Causal copilot: Robust causal analyses at scale

End-to-end pipeline

- Causal Roadmap

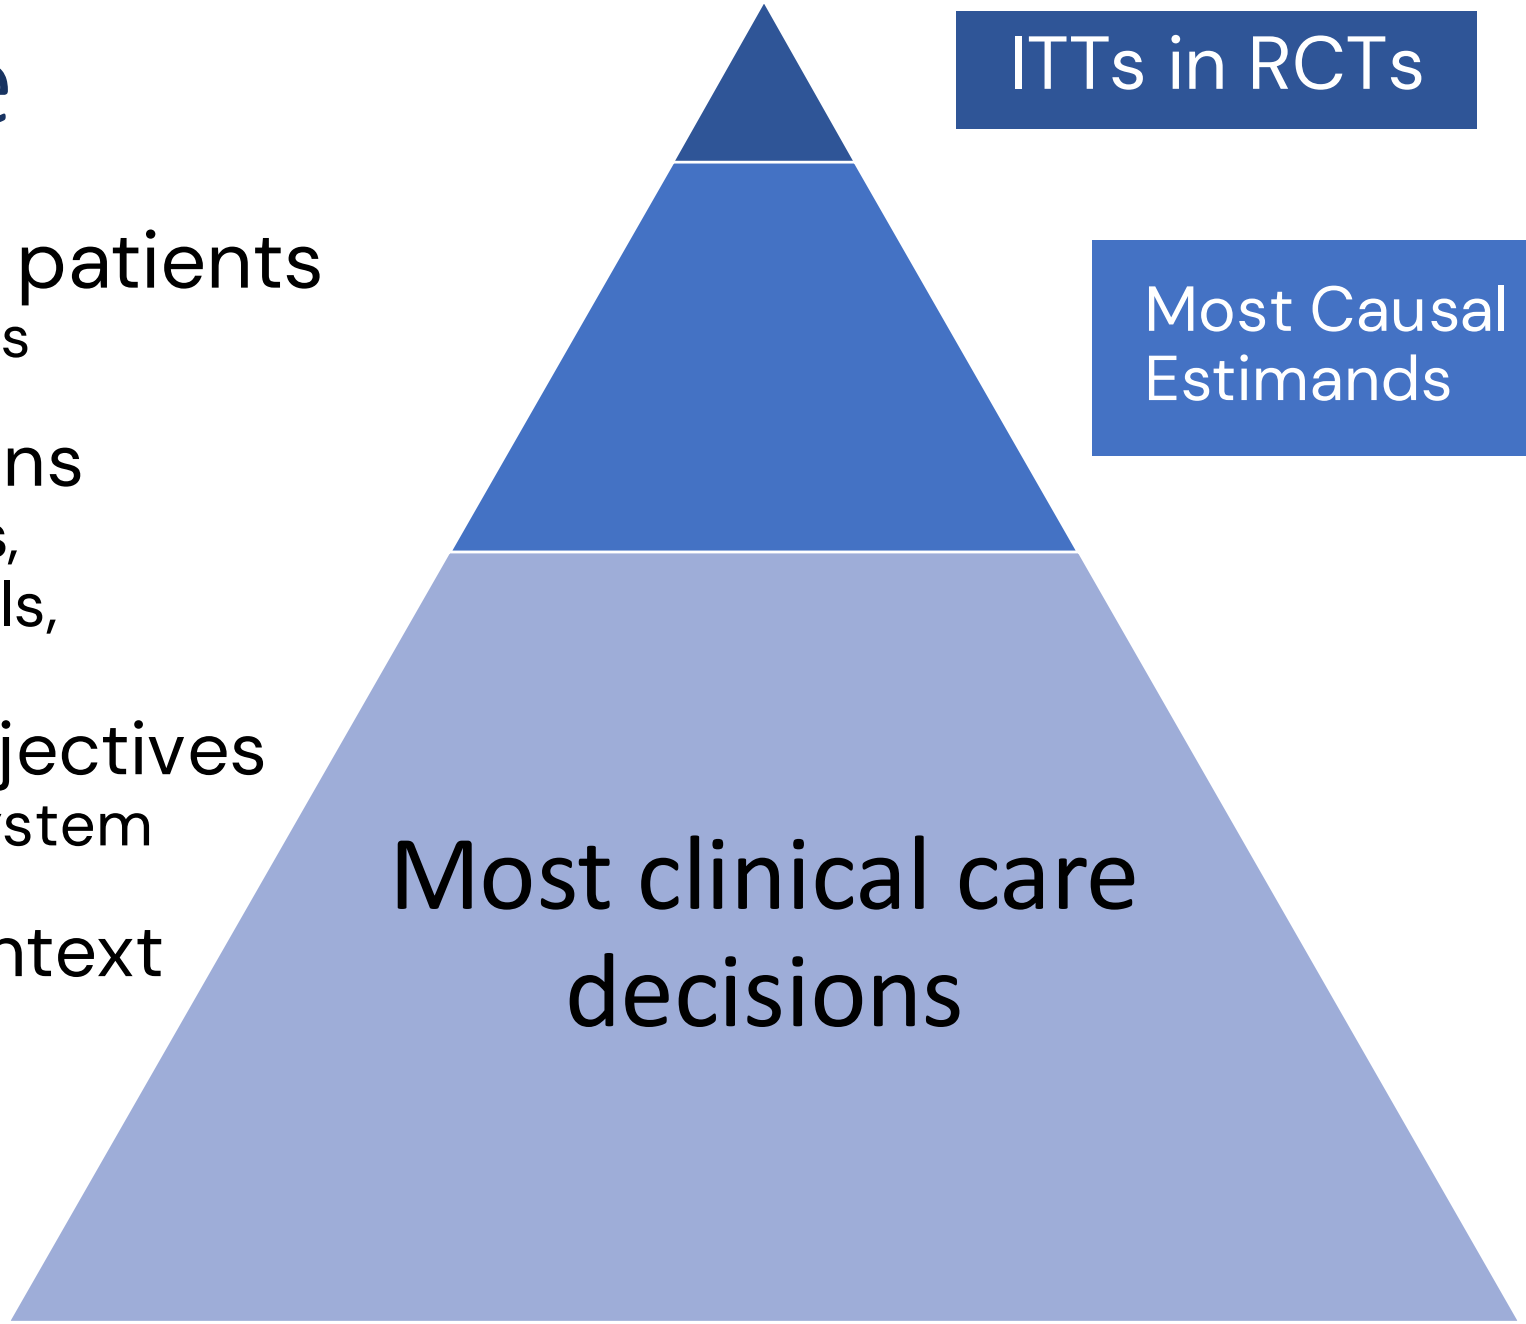
Agentic Harness

Connect frontline users with fast robust cutting-edge causal toolkit



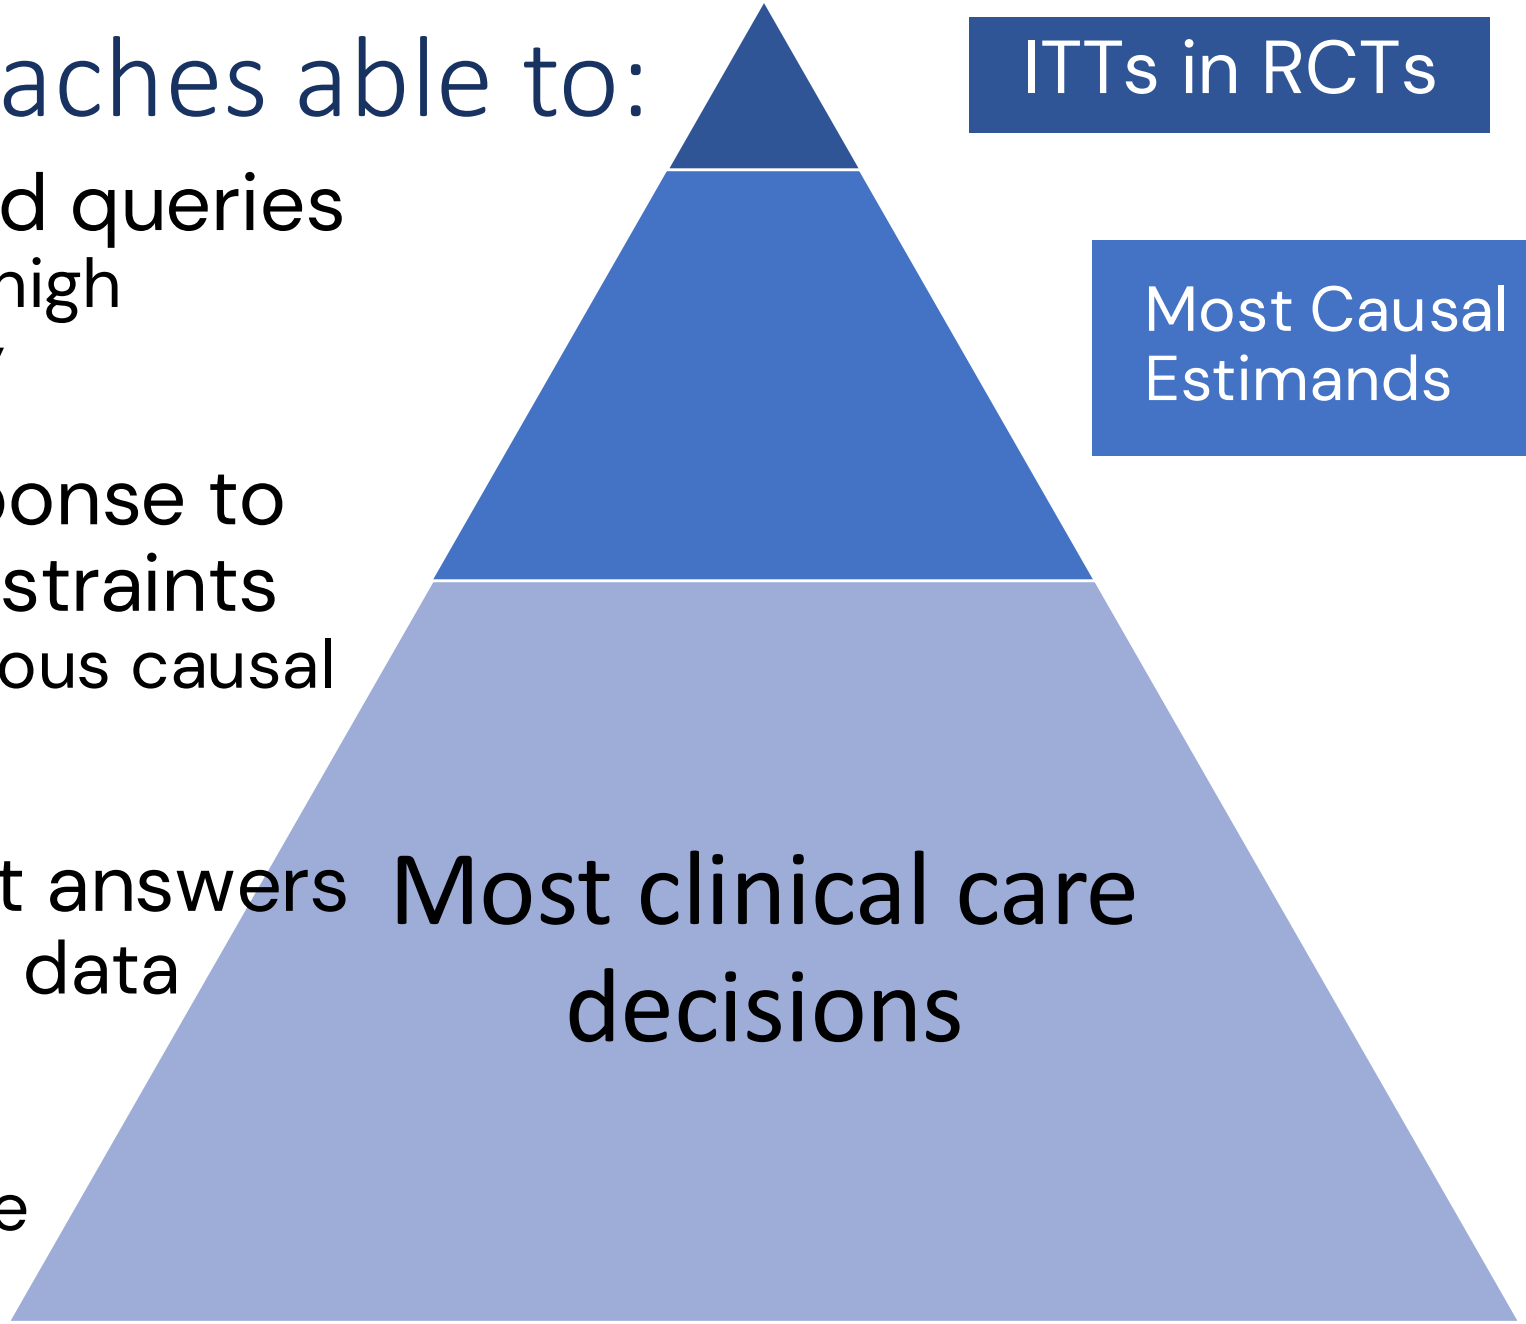
Most clinical care

- Complex individual patients
  - Ex. multiple conditions
- MANY possible actions
  - Labs, Imaging, Meds, Procedures, Referrals,
- Multiple dynamic objectives
  - Patient, provider, system
- Variable evolving context and constraints

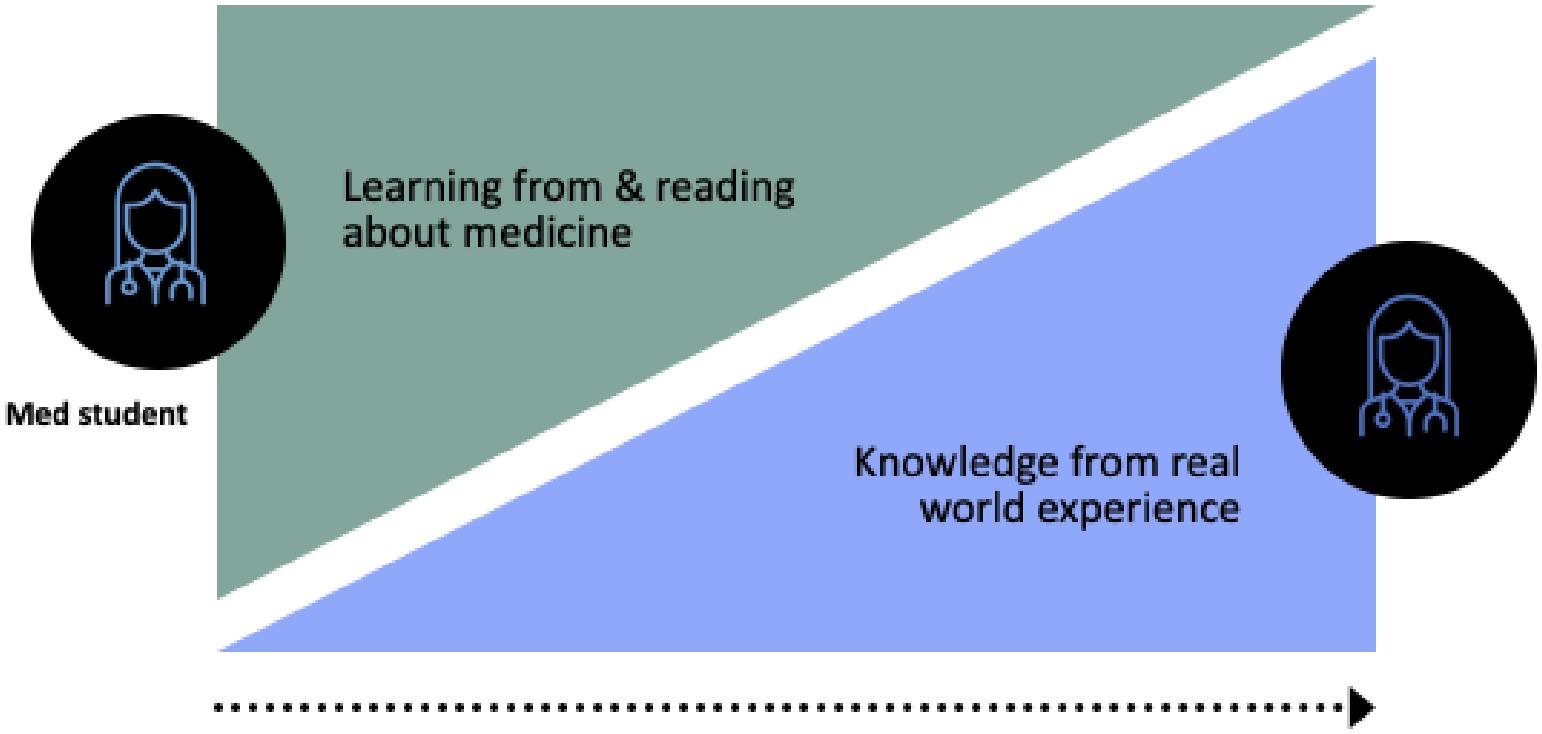


Need causal approaches able to:

- Answer individualized queries
  - Inputs: multimodal high dimensional history
- Modify these in response to local objectives/constraints
  - For complex ambitious causal estimands
- Generate fast robust answers
  - Using limited local data
- Adapt rapidly
  - Over time and place



Clinical AI must understand how clinicians and systems practice medicine



Most clinical knowledge is not written down.

It is learned through hundreds of hours of experience treating real world patients in the healthcare system.

Knit Health

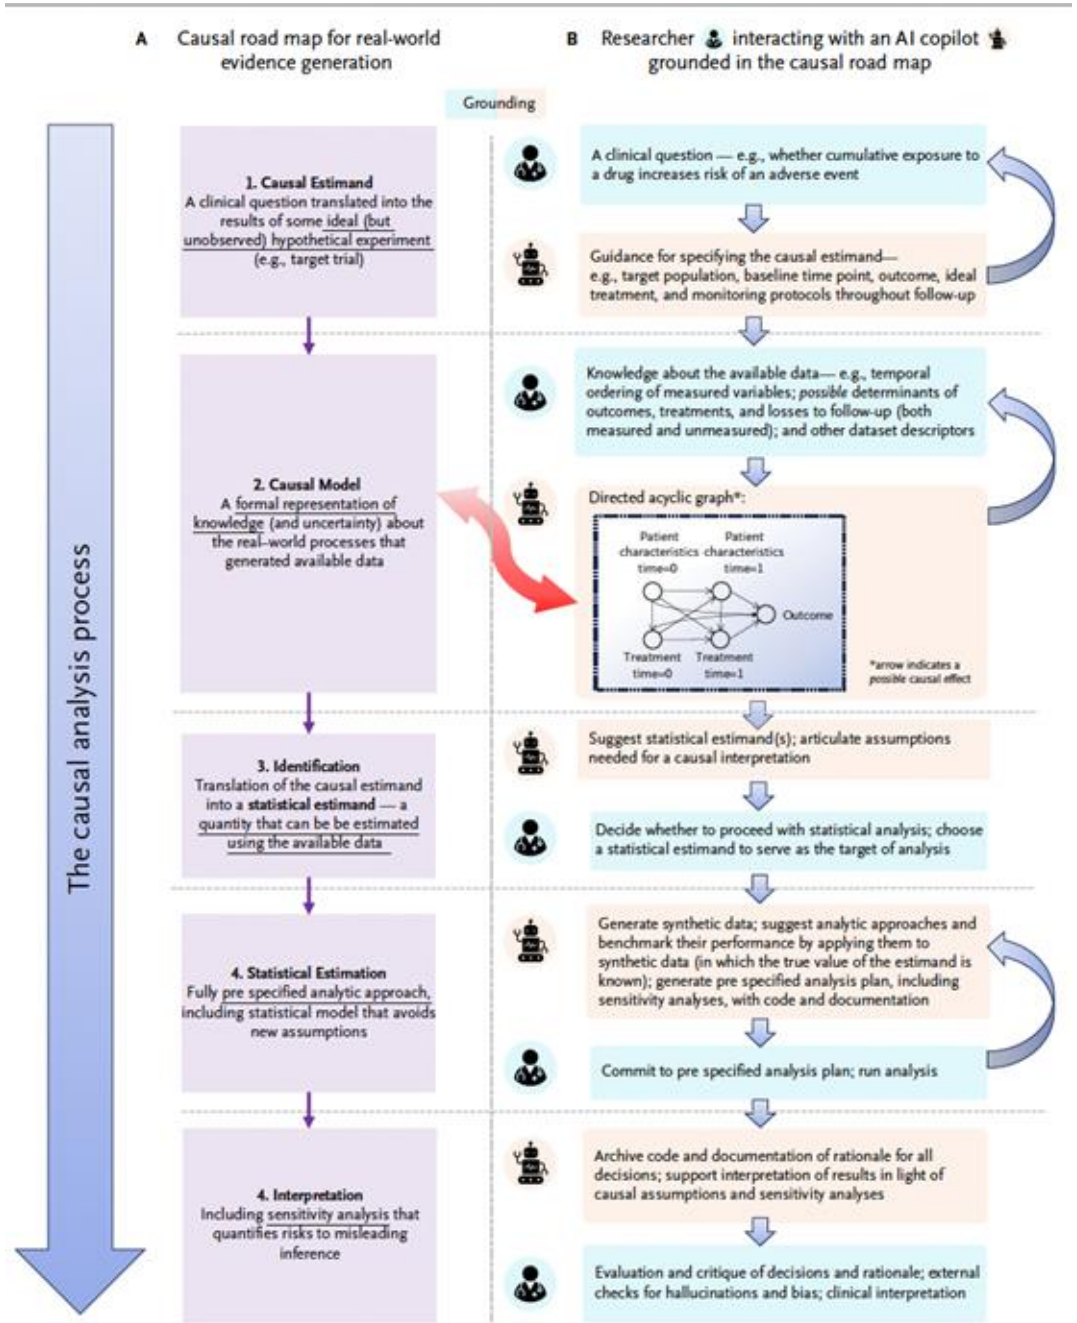
Integration of two approaches

End-to-end pipeline

- Causal Roadmap

Connect frontline users with fast robust cutting-edge causal toolkit

Multimodal foundation model as input to “causal ML” estimators



Thank You!