The Intersection of Behavioral Economics and AI

August 5, 2026

Jonathan Kolstad
UC Berkeley and Knit Health

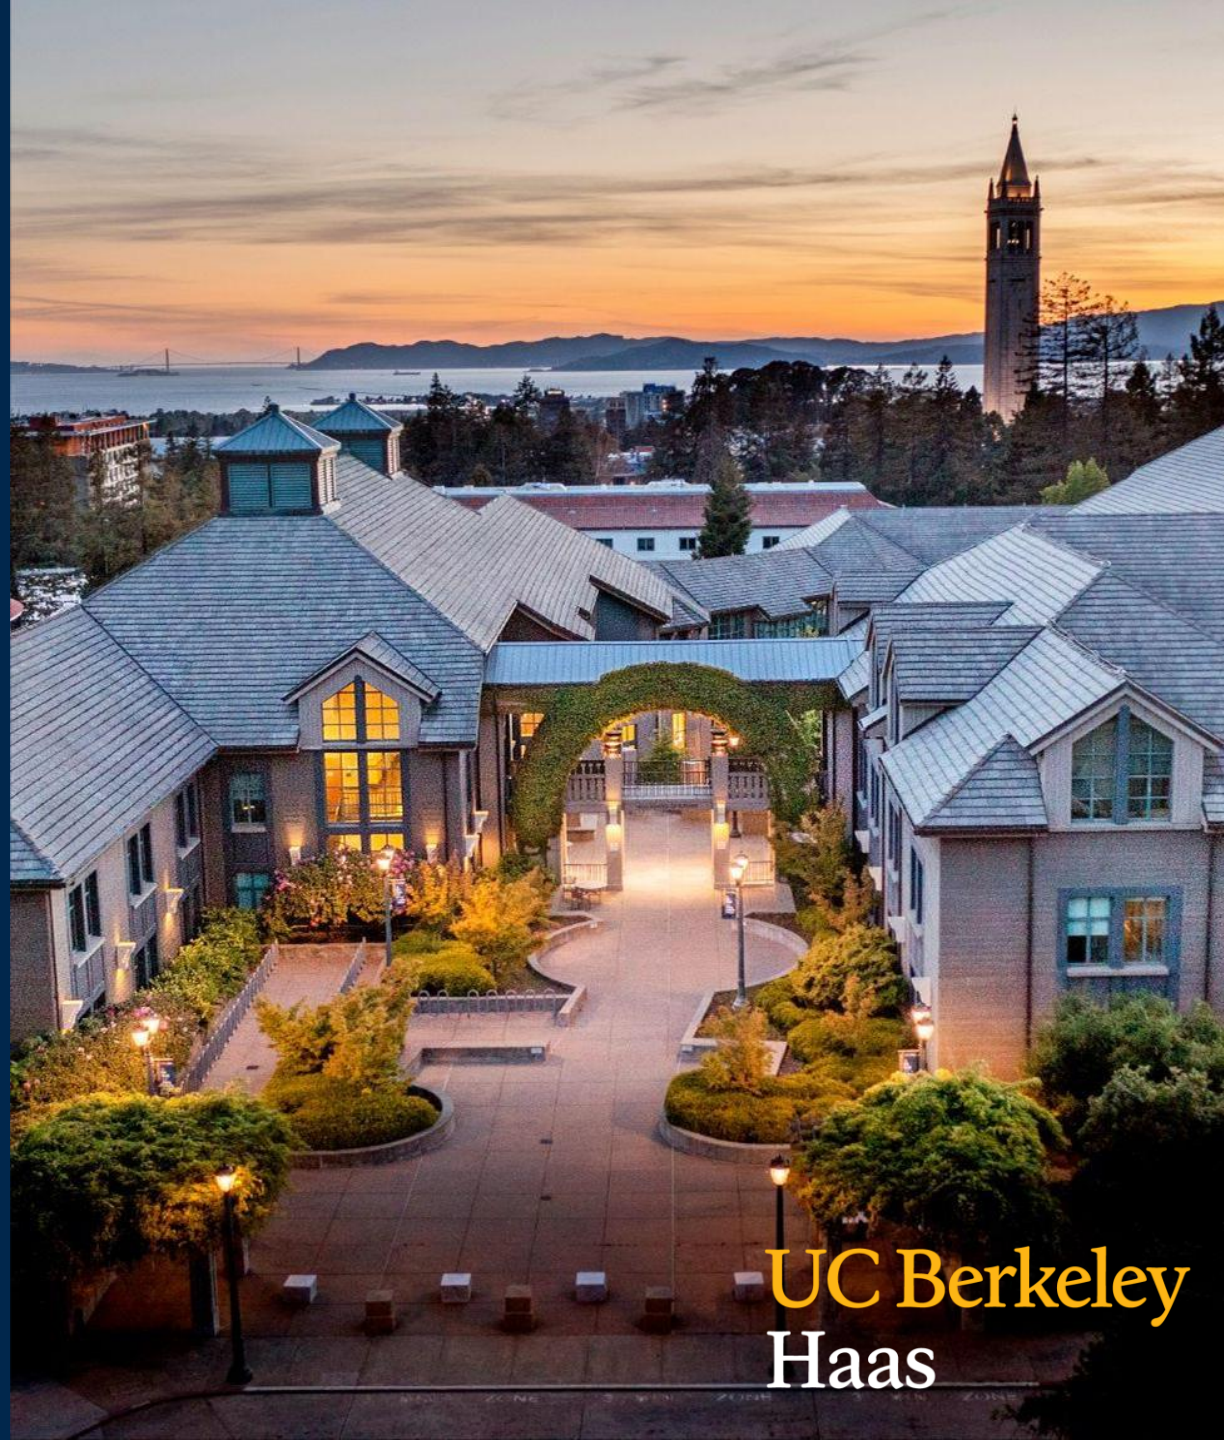

Introduction and Agenda

- Healthcare decisions and the application of AI
- Behavioral economics and cognitive load in healthcare
- World model of health care (Large Clinical Behavioral Model)
- Creating AI that is built for healthcare (not just language)

Who am I?

- Economist, focused largely on healthcare
- AI/ML+behavioral economics
- Founder of Picwell, Healthpilot, OODA and, most recently, Knit Health

The Healthcare AI Paradox

- Artificial intelligence is advancing rapidly.
- Across industries we see:
  - autonomous vehicles
  - AI software engineers
  - AI copilots for knowledge work
- Yet healthcare systems still struggle with:
  - access to specialists
  - clinic backlogs
  - physician burnout
  - coordination across care teams

Why hasn’t AI meaningfully improved care delivery/productivity?

Industry Conventional Wisdom

- Most AI development assumes the core challenge in healthcare is:
  - information processing
- Examples:
  - reading medical records
  - summarizing notes
  - predicting risk
  - answering clinical questions

These are important capabilities, but…

- they are not the central constraint.

A Different View

Healthcare is not primarily an information problem.

It is a decision and coordination problem.

Health systems succeed or fail based on how well they:

- make decisions
- route patients
- allocate clinical expertise
- coordinate sequential care decisions

AI becomes transformative when it improves how these decisions are made.

AI that will transform healthcare needs to understand the realities of decision making and the healthcare system

Core Thesis

- Healthcare is not primarily an information problem.
- It is a decision and coordination problem.
- AI will matter most when it improves how decisions are made and routed.

- We need to understand both AI and the behavioral economics of the system

How decisions actually happen in healthcare

UC Berkeley
Haas

Healthcare Is a Sequential Decision System

- Healthcare performance depends on millions of decisions:
  - who gets seen
  - when they are seen
  - which specialist they see
  - what treatment path follows
- Care involves chains of decisions:

symptom → PCP visit → referral → diagnostic testing → treatment

- Small early decisions shape the entire pathway.

Patient Example

- Maria, 58
- Symptoms:
  - persistent fatigue
  - shortness of breath
  - mild swelling
- Her primary care physician suspects possible heart failure.
- What happens next?

The Real Decision

- Maria’s care pathway might involve:
  - PCP evaluation
  - referral to cardiology
  - diagnostic imaging
  - medication decisions
  - follow-up management
- Each step depends on prior decisions.

Cognitive Load and Decision Making

- Maria's PCP is making this decision:
  - during a full clinic schedule
  - with incomplete data
  - with competing diagnoses
- Clinicians rely on experience and heuristics.

Cognitive Load and Decision Making

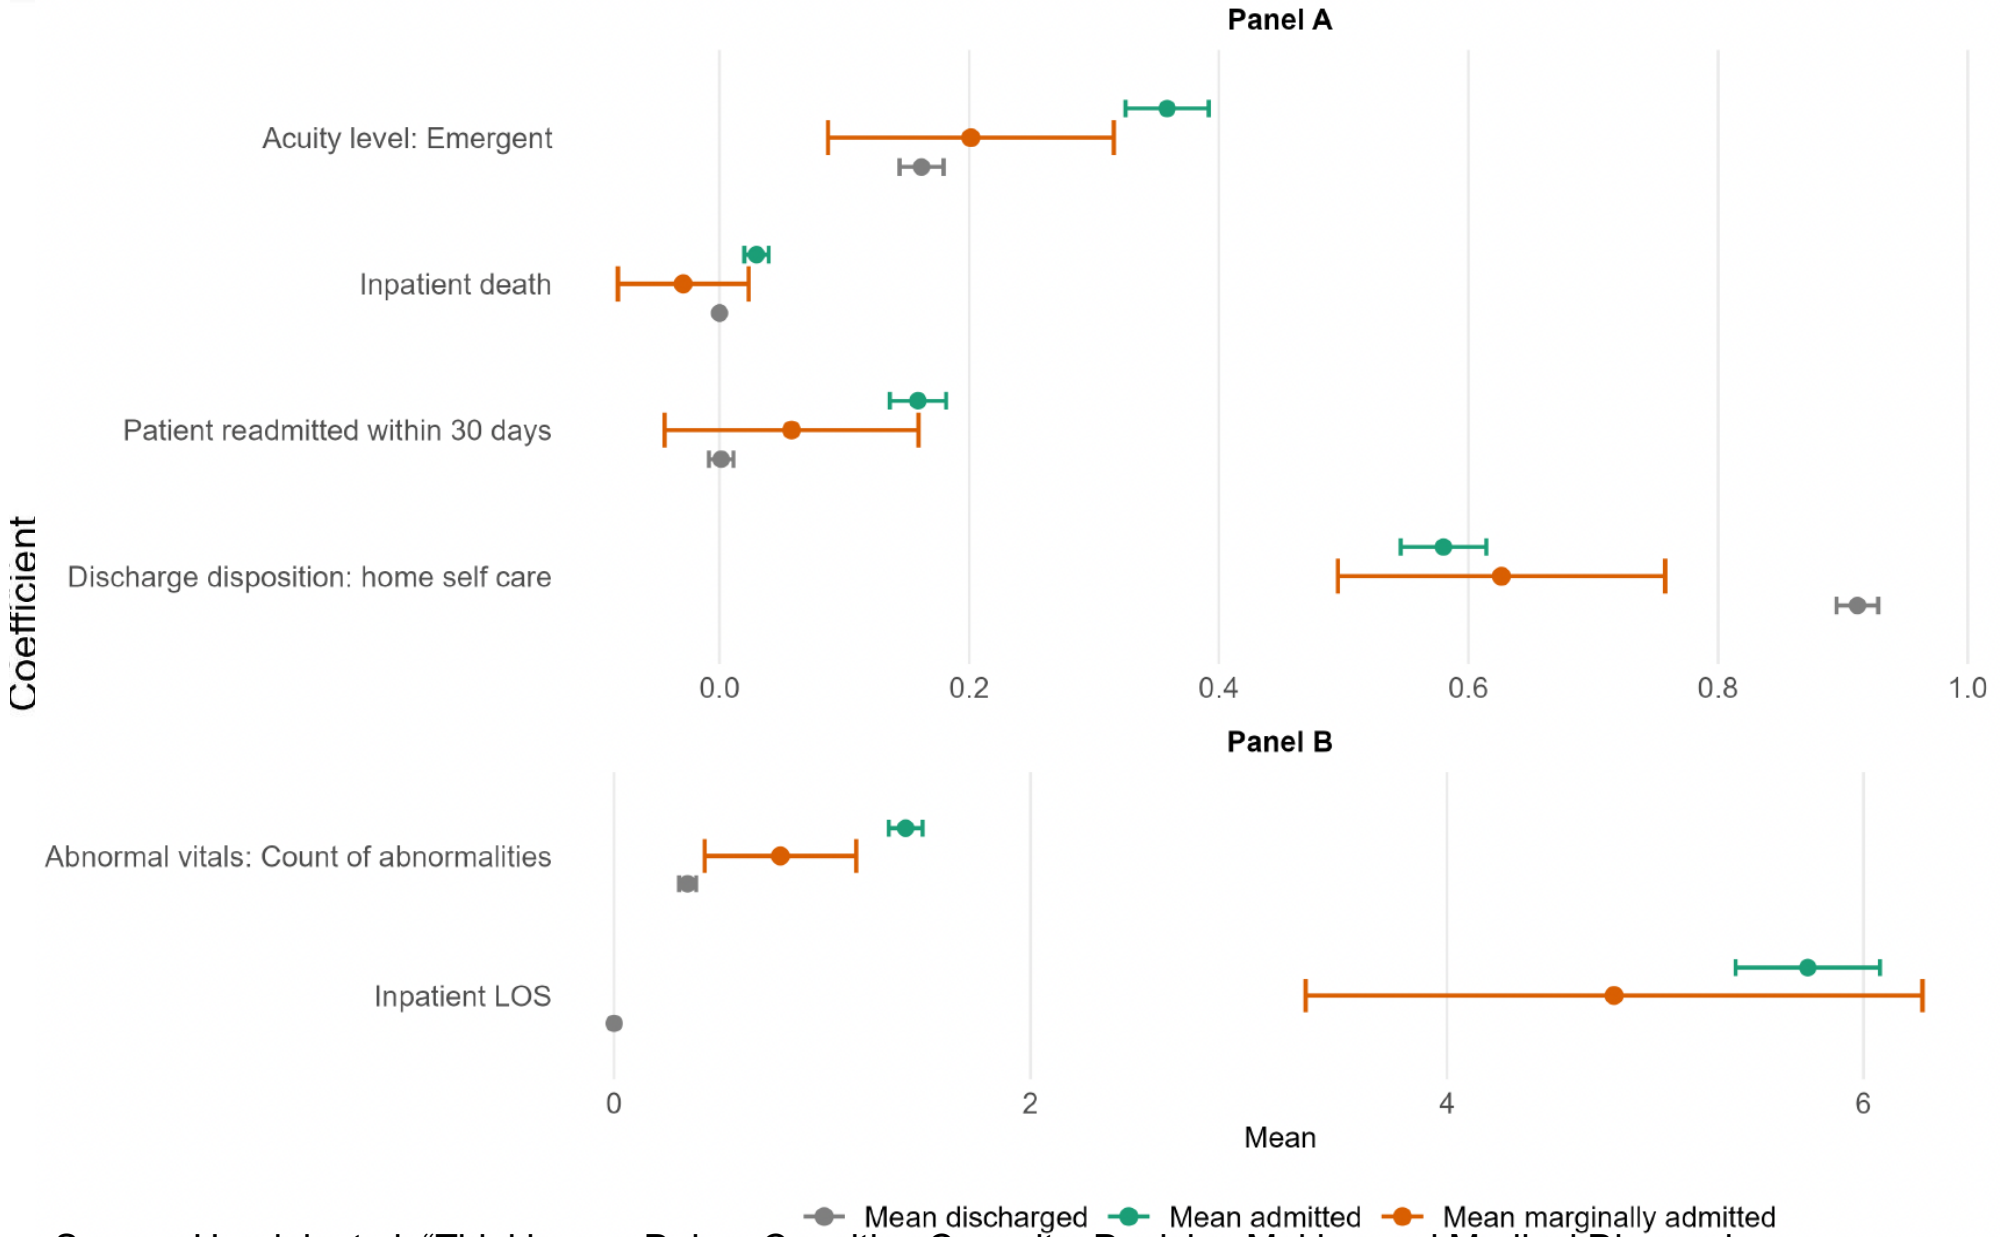

Source: Handel, et al. “Thinking vs. Doing: Cognitive Capacity, Decision Making and Medical Diagnosis

The Real Bottleneck

- Cardiology availability is limited.
- So decisions become:
  - urgent referral?
  - routine referral?
  - treat and monitor?
- Small decisions determine how quickly Maria receives care.

The Current Wave of AI

UC Berkeley
Haas

Predictions vs. Decisions

- An AI model might predict:
  - Maria looks like she could have a heart failure hospitalization.
- But the system still must decide:
  - Should she see cardiology immediately?
  - What test should be ordered?
  - Who should manage her care?

- Much of the knowledge needed to answer these questions is tacitly learned by clinicians and depends on teams and system

Limits of Current AI

UC Berkeley
Haas

Large Language Models

- Language models help with:
  - summarizing records
  - drafting documentation
  - answering clinical questions
- These tools improve information flow.
- Predicated on individual level language tasks as the main form of intelligence

But They Don’t Model Care Delivery

- Language models do not learn:
  - referral patterns
  - treatment decisions
  - specialist allocation
- They model text, not healthcare.

Opportunity Example: A Different Approach to AI

UC Berkeley
Haas

Learn from Clinical Decisions

- What if AI learned from:
  - how clinicians refer
  - how they diagnose
  - how they escalate care
  - how they allocate expertise
- Instead of learning only from text or risk labels.

A World Model for Healthcare: Large Clinical Behavioral Model

- Large Clinical Behavioral Models learn from:
  - millions of clinical decisions
  - pathways through the healthcare system
  - how expertise is applied
- This captures how clinicians actually solve problems.

What is an order path?

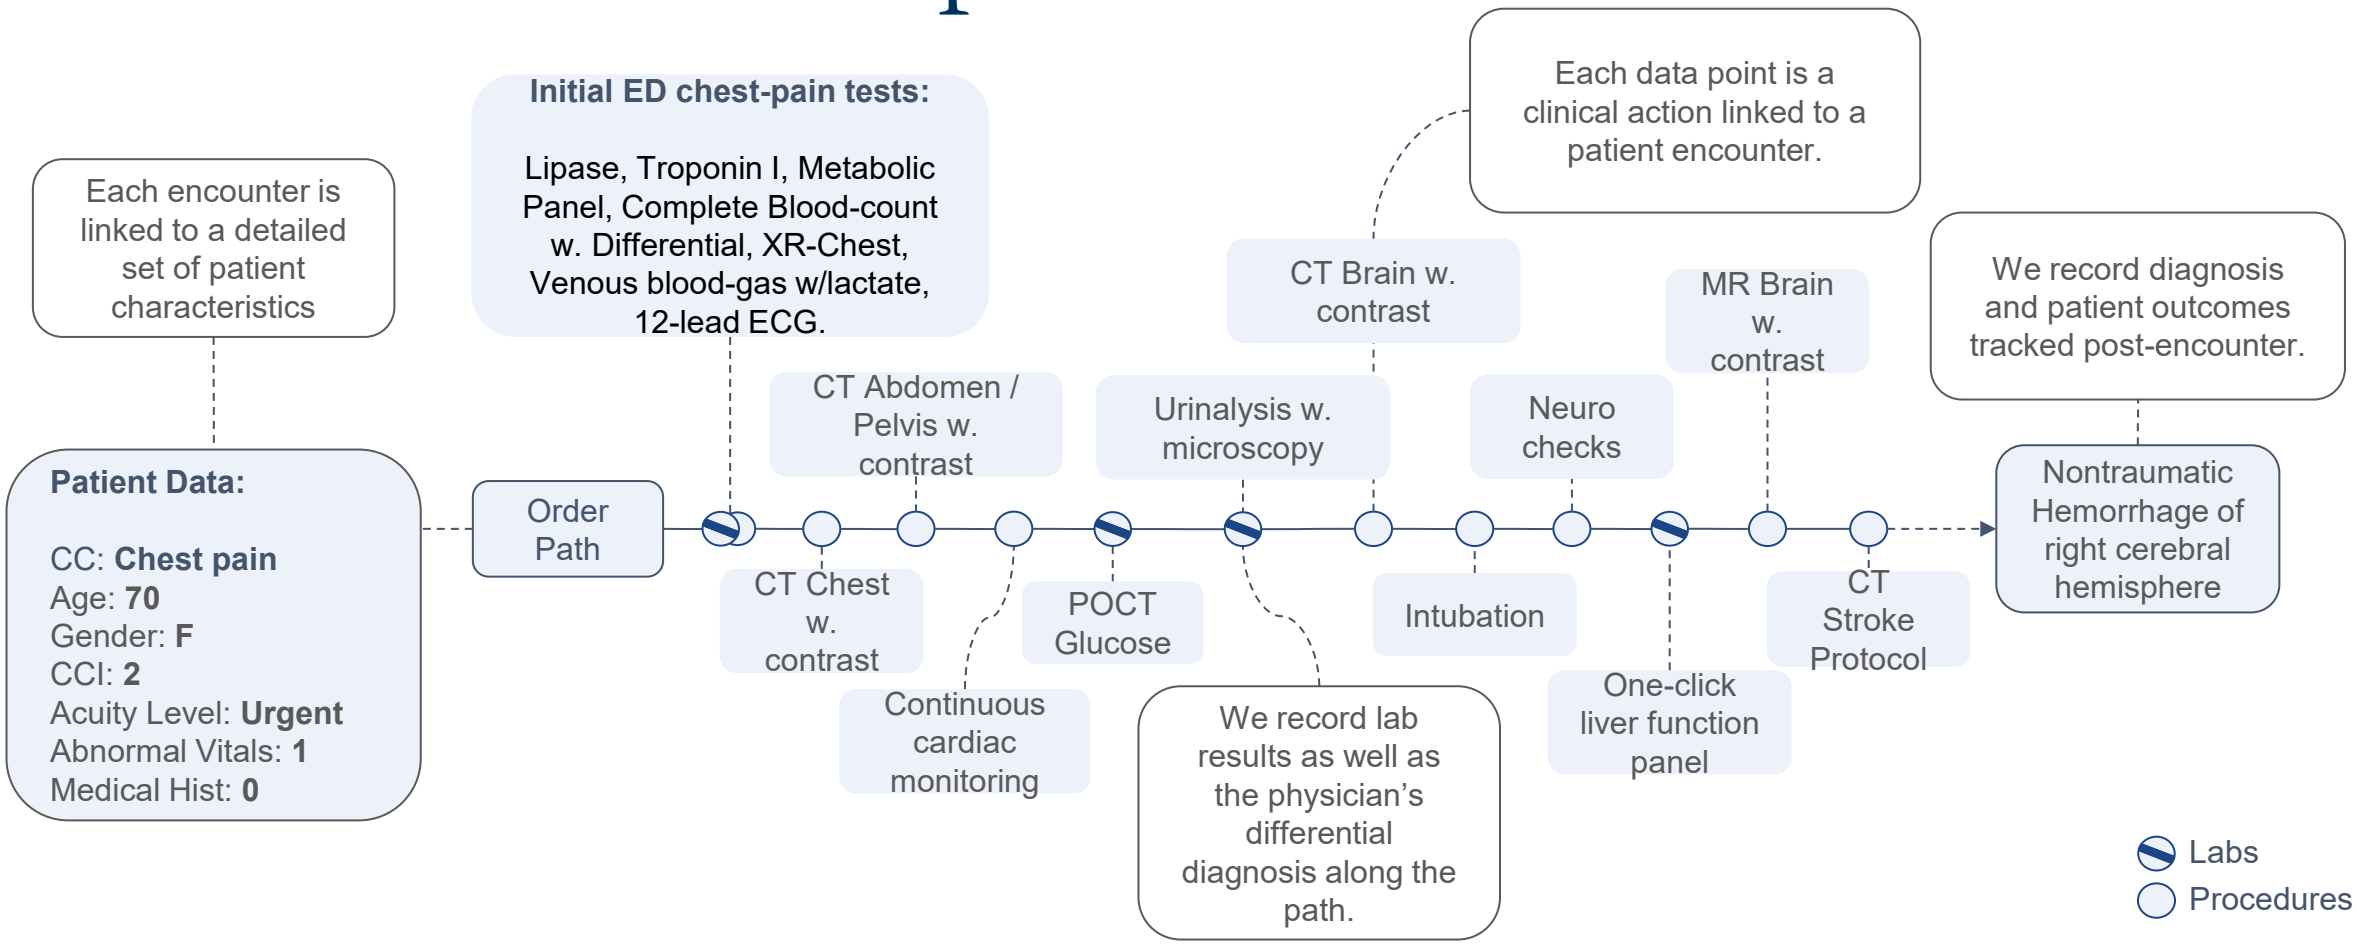

Sample Patient Path

Berkeley | UCSF

Actions and Information

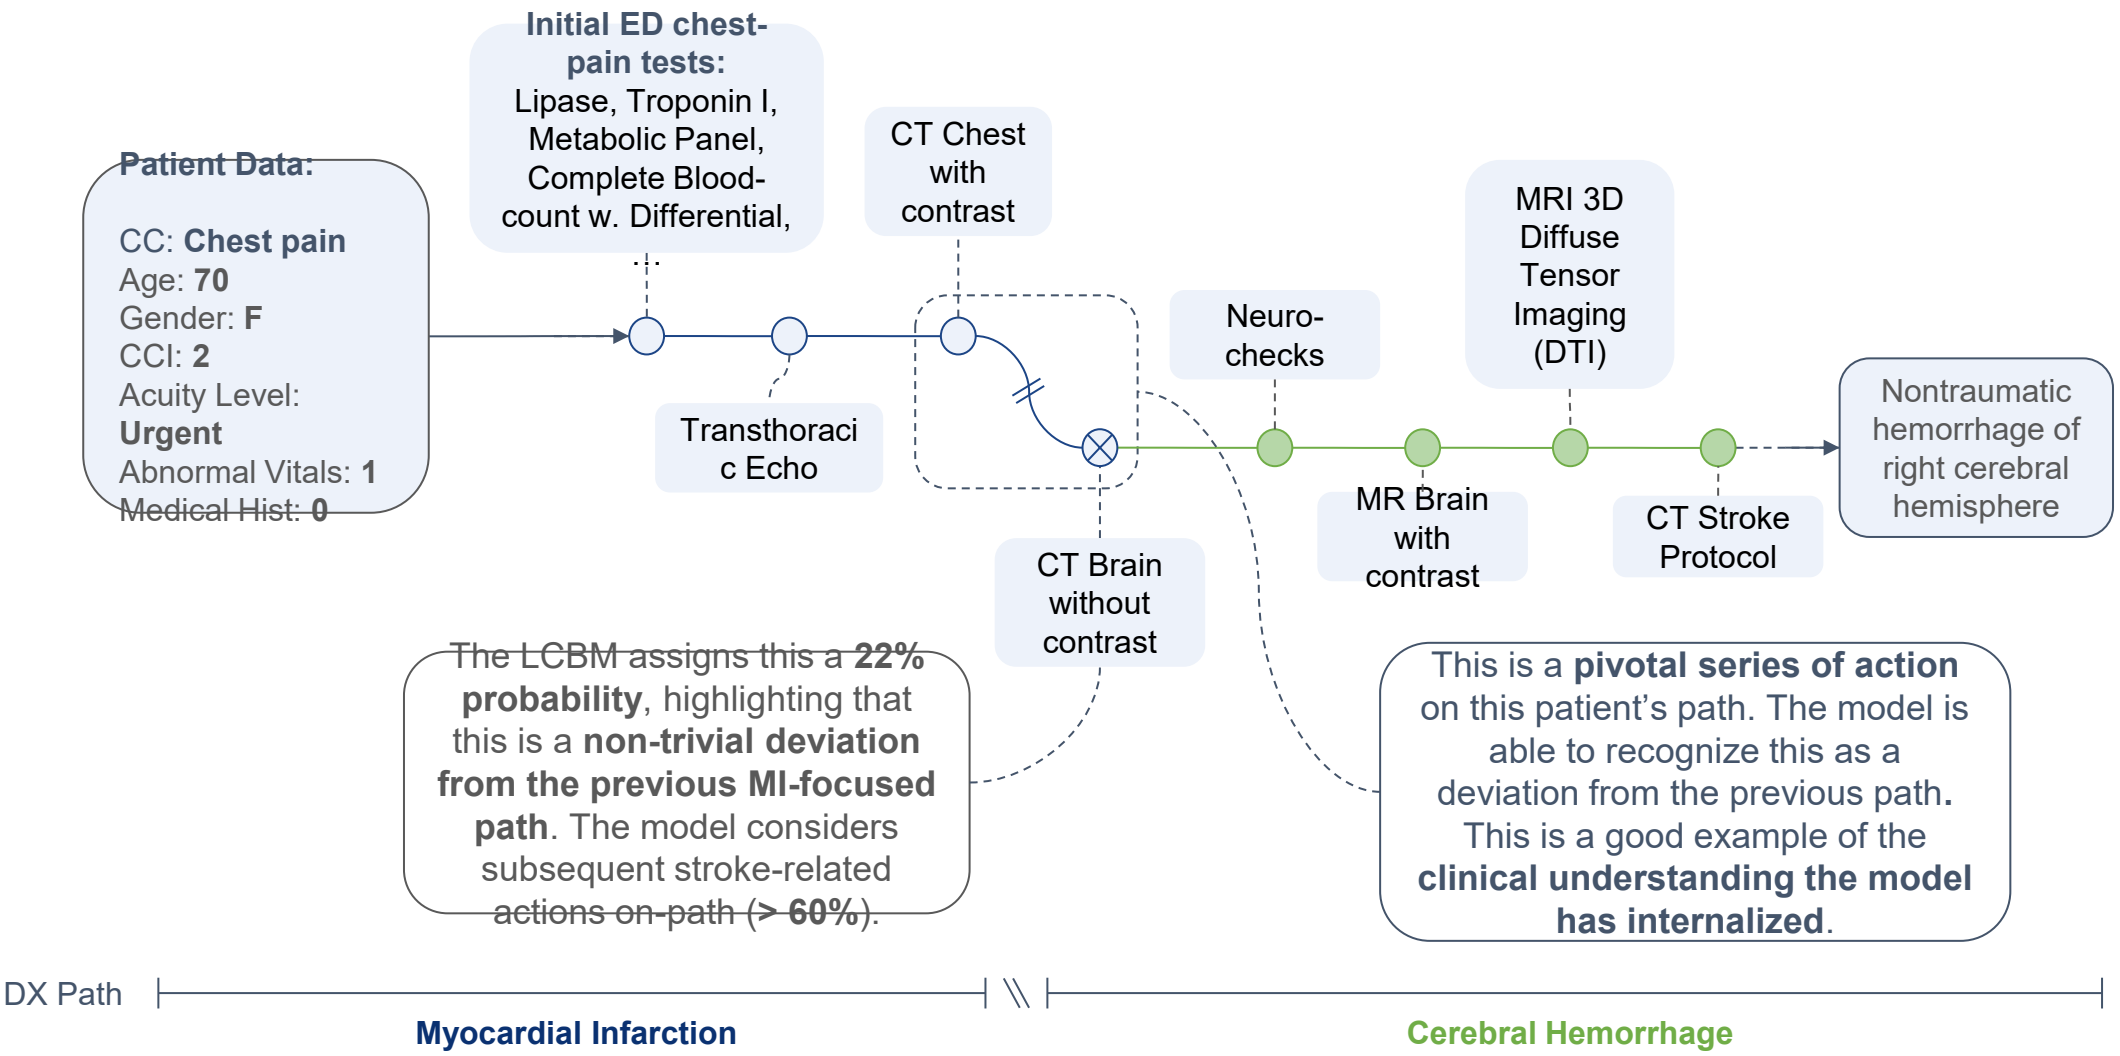

Berkeley | UCSF

The Value of an LCBM for AI in Healthcare

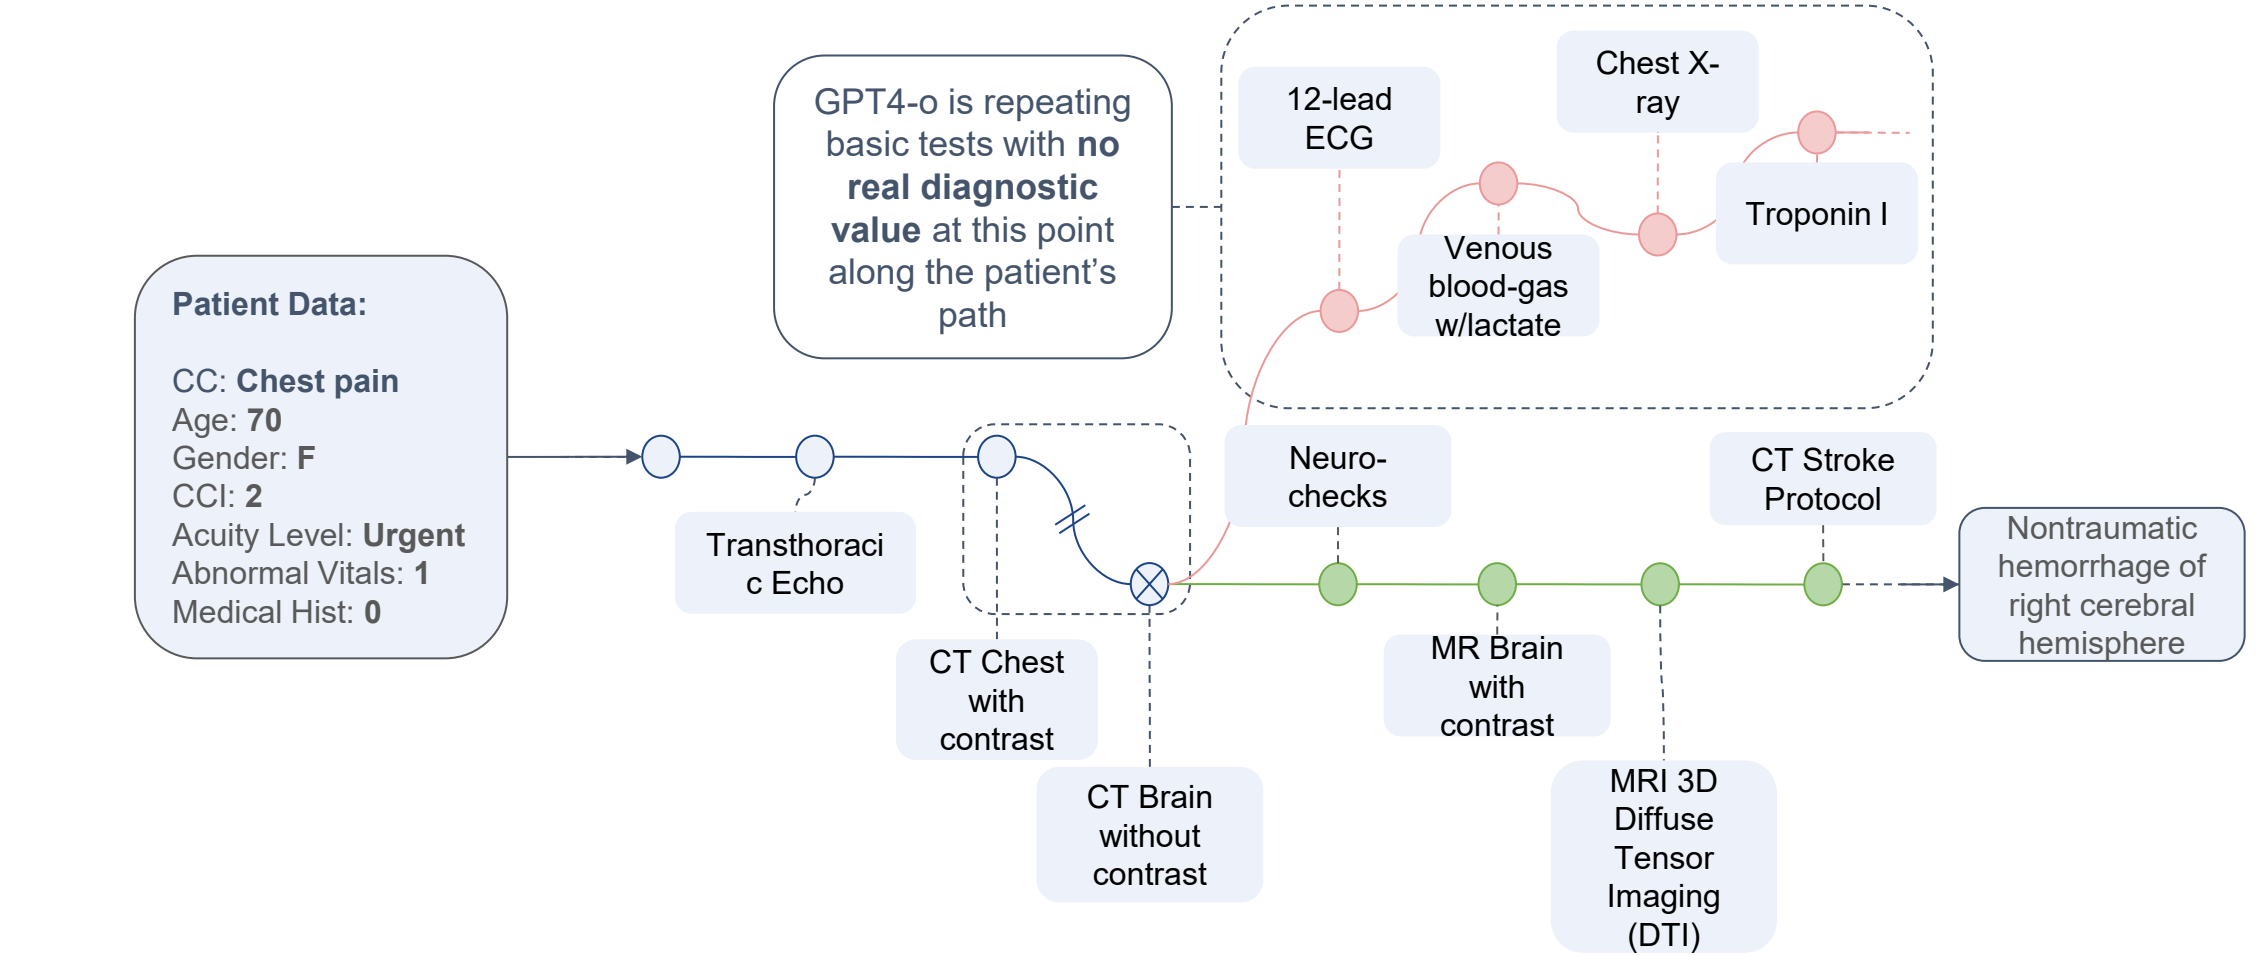

Berkeley | UCSF

Med-BEAM's inputs and outputs are designed for programmatic scale

Med-BEAM natively ingests a patient's history and predicts likely future clinical events. Rolling predictions forward simulates many possible future paths, to each of which we can assign a probability. In this way, Med-BEAM becomes a **versatile prediction platform.** These predictions are delivered programmatically via custom API into operational workflows.

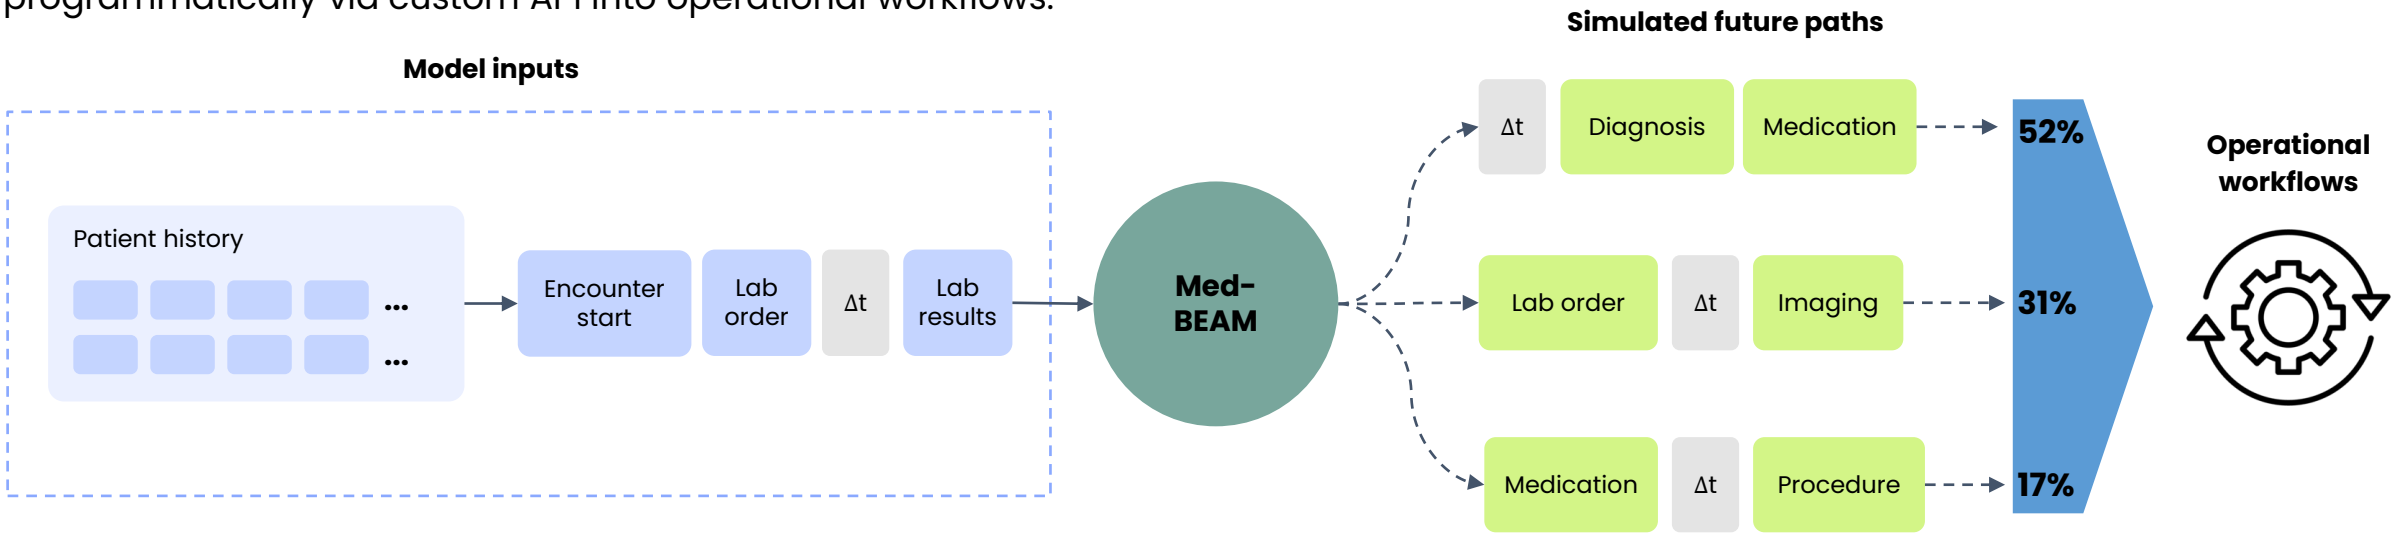

The event vocabulary spans encounters, medications, lab orders and results, imaging, diagnoses, procedures, and the time elapsed between events (Δt).

BerkeleyHaas

Back to Maria

- Now imagine a system that learns from thousands of cardiology referrals.
- It could identify:
  - patients like Maria
  - optimal referral timing
  - appropriate diagnostic pathways
- This supports clinicians rather than replacing them.

World Model Deployment

Self-supervised world model

Supervised: one model per task

Define a prediction task
e.g. 30-day readmission

↓

Label outcomes by hand
chart review, months of clinician time

↓

Fit a task-specific model
one model, one outcome

↓

Repeat for the next task
× every outcome you care about

LLMs: the wrong manifold

Trained on language
notes, textbooks, papers: descriptions of care

↓

Native tokens: words
real trajectories live outside the language distribution

↓

Can describe care, not predict it
no calibrated probabilities over the next clinical event

Med-BEAM: native to care

Raw EHR event stream
millions of patients, billions of events, no annotation

↓

Predict the next event
the event that actually happened is the label: supervision for free

↓

One world model
native tokens are the primitives of care: a prediction platform by construction

…and that's it. New tasks reuse the same model.

BerkeleyHaas

Benchmarking

A clinical world model does what LLMs cannot

2.2–5.5×
more clinically accurate

33–73×
faster per query

17–399×
cheaper per query

2.6–38×
more calibrated: probabilities reflect reality

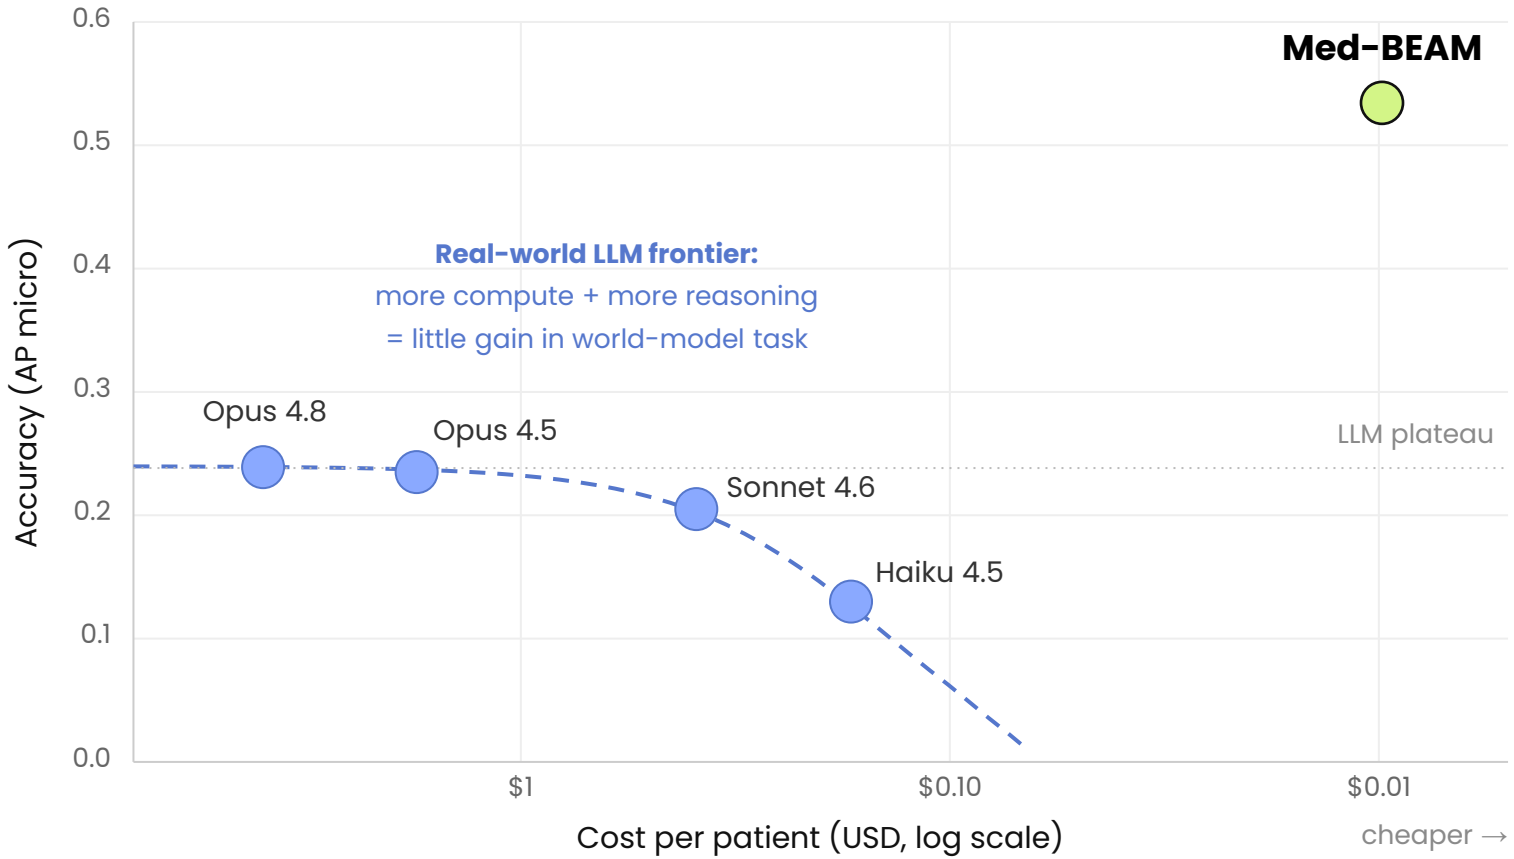

BerkeleyHaas

Zooming Out: Lessons from Other Industries

UC Berkeley
Haas

AI Works when Embedded in Decisions

- Industries that saw major gains:
  - Aviation
  - Logistics
  - Finance
- AI improves systems when it supports operational decisions, not just analysis.

Leadership Implications

- Successful organizations will:
  1. build decision infrastructure
  2. integrate AI into clinical workflows
  3. invest in data quality
  4. focus on operational outcomes
- Not just model performance.

Conclusion: The Opportunity in Healthcare

- Healthcare generates enormous collective clinical intelligence.
- AI can help systems learn from that intelligence.
- The goal is not replacing clinicians.
- It is helping patients like Maria reach the right expertise at the right time.
- Requires understanding realities of system and human behavior

Thank you!
Feel free to reach out

jkolstad@Berkeley.edu; jkolstad@knit.health