

Agentic AI in Healthcare

1 September 2025

Introduction to Agentic AI and its role in Healthcare

Samik Ghosh, COO, SBX Corporation

Samik Ghosh's career spans science, engineering and entrepreneurship. He serves as COO at SBX Corporation Tokyo group, CEO of SBX Technologies, a SBX subsidiary company focused on healthcare and is a Senior Scientist at The Systems Biology Institute, Tokyo.

He holds a PhD in Computer Sciences from University of Texas and has diverse experience in life-science and healthcare industries in India, USA and Japan.



Preserving Agency in Healthcare

Agentic AI and its role in healthcare

Samik Ghosh, PhD
ghosh@sbx-corp.com

Agentic AI in Healthcare Workshop

29 August

2025 A.D

A.D = Agent Dominance

Journey began 25 years ago in “systems thinking”

25 years of Innovation


- ◆ Over 2 decades in the business of Biology
- ◆ Category creator in “**Systems Biology**”
- ◆ Pioneer at the intersection of **biology, data & compute**
- ◆ Bootstrapped, profitable and stable organic growth over a decade

The Systems Biology Institute (SBI), Tokyo, Japan is a **non-profit private research institution established in 2000** to promote systems biology research and its application to medicine and global sustainability. SBI focuses on the rapid translation of basic research to practical outcomes for both business and clinical applications. SBI has been engaged in a series of research programs funded by government grants, both national and international as well as extensive private contracts for industrial applications of systems biology research.

SBX Corporation, Japan (SBX Group and its subsidiaries) is a **tech.bio spin-off from SBI, Tokyo, established in 2011**, focused on digital biology products and services to enable drug discovery and healthcare. Driven by deep science, powerful platforms, and precision engineering, the diverse team of SBX spread across Japan, North America, and India harnesses the power of data and intelligence to provide bespoke solutions for global companies in drug discovery, health, and wellbeing.

Global Talent Network

Locations



Building agents for craftsmanship



Connectivity and automation platform

Powered by Garuda™ is a technology framework, enabled by the **Garuda platform**, which provides connectivity and automation of diverse data, devices and analytics to create novel pipelines, products and solutions in various science and industry verticals in platform-agnostic manner.



Text analytics and intelligence platform


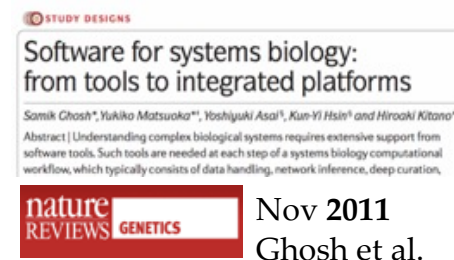
Taxila is a living system which grows and learns by forming new connections from an ever-accumulated knowledge.

Powered by Taxila™ is a technology framework, enabled by the **Taxila platform**, which provides automatic context-aware aggregation and search of relevant information, driven by AI-powered mining and analytics engine for driving actionable insights with an intuitive user interface to create new and own insights on the topic and resource of your interest.



Biomedicine focused intelligence platform

Gandhara provides Wisdom of the Crowd AI approach which integrates a large number of machine learning and deep learning algorithms for big data analytics.



The Systems Biology Institute (SBI) participated in the AstraZeneca-Sanger Drug Combination Prediction DREAM Challenge, and ranked 4th amongst 142 global submissions

The AstraZeneca-Sanger Drug Combination Prediction DREAM Challenge
(Registration Open: 25 June 2015, Close: 21 March 2016)
<https://www.synapse.org/#!Synapse:syn4231880/wiki/>

Nat Commun. 2019 Jun 17;10(1):2674.
<https://www.nature.com/articles/s41467-019-09799-2>



Shades of AI

Rule-based, rigid and template-driven. Focus on automation (RPA) for efficiency.
Example: Rule-based report generation and checking

Systems that employ learning models on static datasets and solve specific problems
Example: Supervised learning for classification of documents,

Systems which are trained on static datasets for specific problem but with capability to augment their learning from new data in dynamic manner.

Generative AI systems

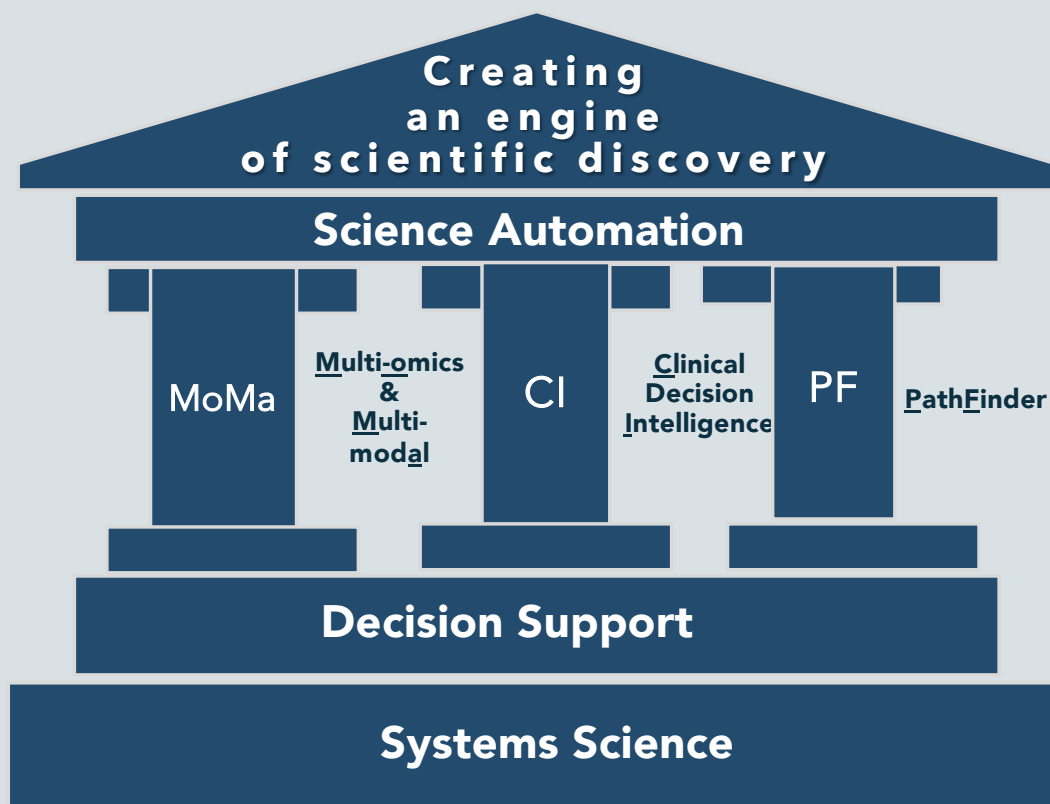
Pre-trained large-scale models (language models) trained on vast datasets with the ability to generate output based on specific prompts

Abstract

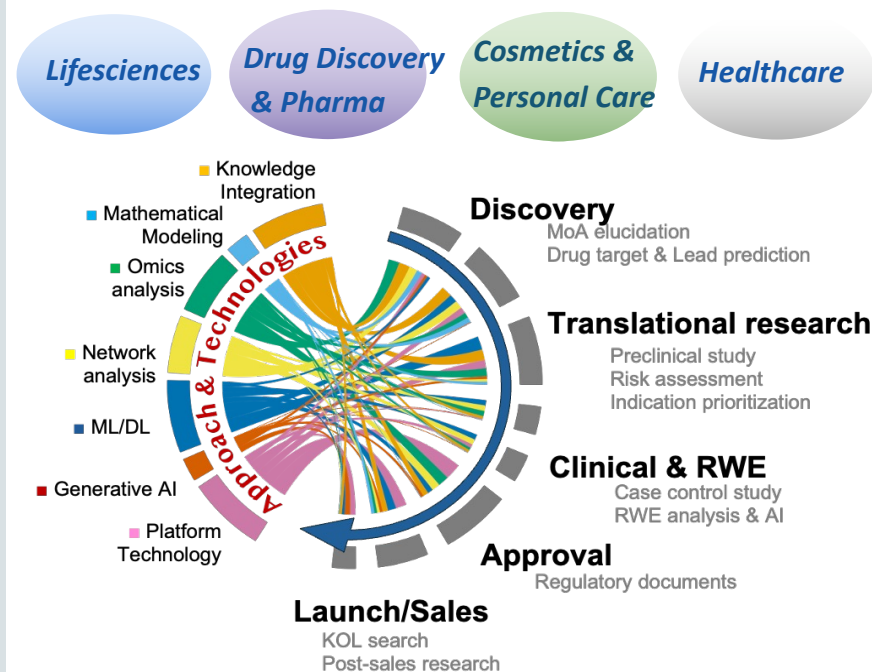
Over the last few years Literature-based Discovery (LBD) has regained popularity as a means to enhance the scientific research process. The resurgent interest has spurred the development of supervised and semi-supervised machine learning models aimed at making previously implicit connections between scientific concepts/entities explicit based on often extensive repositories of published literature. Understanding the temporally evolving interactions between these entities can provide valuable information for predicting the future development of entity relationships. However, existing methods often underutilize the latent information embedded in the temporal aspects of interaction data.

In this context, motivated by applications in the food domain—where we aim to connect nutritional information with health-related benefits—we address the

Bridging the gap from "cure" to "care"



Outcomes *not* Output



Over 35 Partners across across academia, pharma, healthcare and medical technologies

Top 3 Pharma in Japan

Fortune 500 Global companies including consumer health players

Complexity of Care



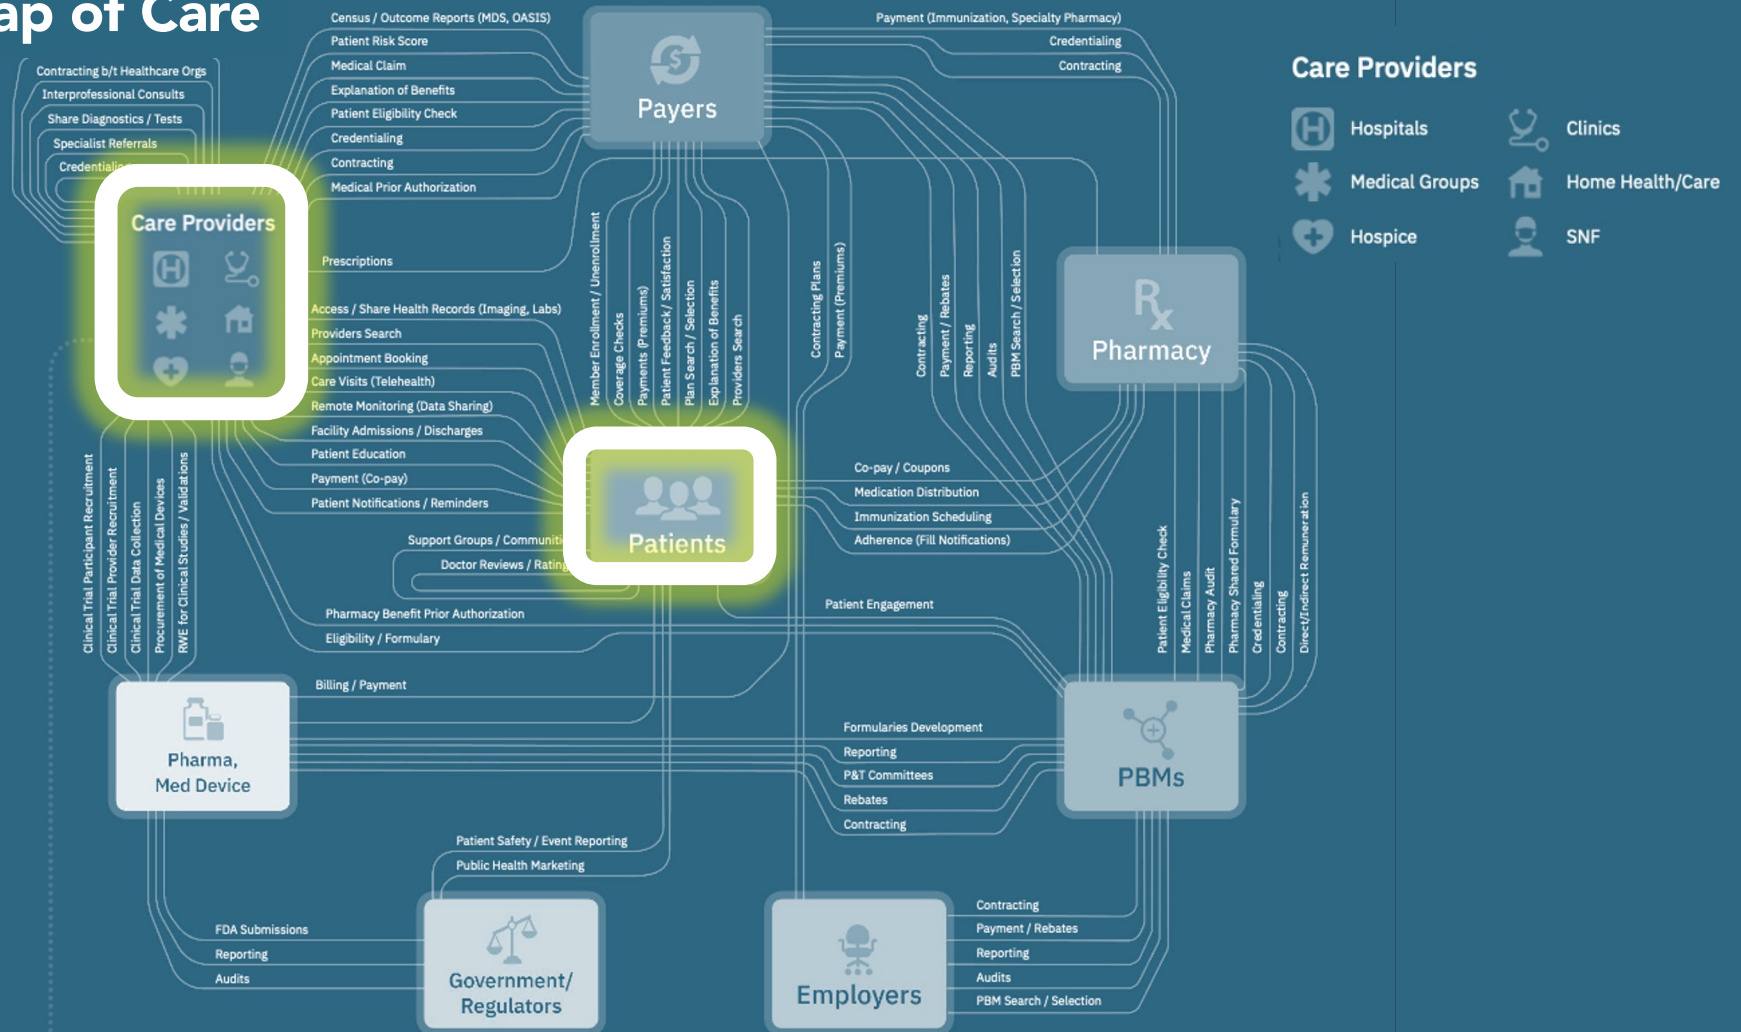
***“The hospital [is] altogether the most
complex human organization ever devised.”***

-- Peter Drucker

Peter F. Drucker. They're Not Employees, They're People, HBR Magazine 2002 Article
<https://hbr.org/2002/02/theyre-not-employees-theyre-people>

We need a "compass" & "map" to navigate

The Subway Map of Care



Care is DUST(y) - **D**ifficult, **U**ncertain, **S**cary & **T**edious

Diagnostic laboratories

Analyze and provide accurate and identifiable data on samples to clinic



Clinic operations & support staff

Research Laboratory/ Pharma trials

- Conduct research on new formulations and drugs
- Perform clinical trials
- Collect and analyze data



Doctors

Serve patients better
Serve more patients



Patients


Get better and continuous advise and consultation on health and well being



Clinic

- Provide better healthcare outcomes to patients
- Improve operational efficiency and productivity
- Increase sales/ value to patients

Medical equipments




Happy Providers make Happy Patients

Death by a thousand “clicks” and “cuts” (of paper)

The Silent Cost of Clicks

In a typical 3-hour clinic session, a significant portion of a doctor's time is consumed by digital administrative tasks. This chart shows the breakdown, highlighting the time lost to clicks, scrolling, and typing instead of patient care.

A Doctor's Time in a 3-Hour Clinic (180 mins)



Digital Admin Work

Patient Care & Conversations



Dr. Zhong Wei Khor · 2nd
NHS Cancer Doctor. We connect He...
[Visit my website](#)
17h · 🌐

I spent 1 hour and 43 minutes in front of my computer during a 3-hour clinic.

Not seeing patients.
Not having conversations.
Not delivering care.

Just clicking, scrolling, and typing.

Out of curiosity, I decided to audit my time—tracking how much of my clinic was spent on a computer rather than with my patients.

I started the stopwatch when I began my clinic, paused it every time I had a patient with me, and restarted it once I was back at my computer.

I spent more time on the computer than with my patients.

So, what was I doing?

- Checking previous clinical letters, scans, and blood results.
- Ordering blood tests and investigations.
- Approving chemotherapy prescriptions.
- Reviewing dictated letters.
- Following up on results and imaging.

Essentially, admin—digital paperwork.

All necessary. None easily replaced.

But it struck me: digital transformation in healthcare has often just turned paperwork into computer work.

We don't need more digital paperwork.


We need better digital solutions—technology that helps us do all this with fewer clicks and less time.

Because every extra click takes away from patient care.

And that's a trade-off we shouldn't have to make.

01 : 43 : 26 . 77


Standing on the threshold of a new Era



System 1 vs System 2 Thinking

- This leap from pre-trained instinctual responses ("System 1") to deeper, deliberate reasoning ("System 2") is the next frontier for AI.
- **It's not enough for models to simply know things—they need to pause, evaluate and reason through decisions in real time.**

<https://www.sequoiacap.com/article/generative-ais-act-o1/>
<https://storage.googleapis.com/deepmind-media/Era-of-Experience%20The%20Era%20of%20Experience%20Paper.pdf>




Welcome to the Era of Experience

David Silver, Richard S. Sutton*

- We stand on the threshold of a new era in artificial intelligence that promises to achieve an unprecedented level of ability.
- A new generation of agents will acquire superhuman capabilities by learning predominantly from experience.

With new opportunities in healthcare

How this AI wave differs from the prior software revolution




To color the white space of operations

Heat Map of Health Care Tasks, Natural Language Processing (NLP) and Natural Language Understanding (NLU) Tasks, and Dimensions of Evaluation Across 519 Studies

Health care tasks							
Enhancing medical knowledge	222	91	44	33	16	10	3
Making diagnoses	100	38	11	11	14	4	0
Educating patients	88	68	32	22	18	3	2
Making treatment recommendations	47	22	9	8	3	1	0
Communicating with patients	35	29	8	15	22	1	0
Care coordination and planning	36	24	4	5	7	1	0
Triaging patients	24	7	5	2	8	8	0
Carrying out a literature review	18	7	2	2	2	2	0
Synthesizing data for research	16	7	2	3	2	2	0
Generating medical reports	8	8	2	0	3	0	0
Conducting medical research	8	7	3	3	3	0	0
Providing asynchronous care	8	5	3	3	1	1	0
Managing clinical knowledge	5	5	1	1	0	0	0
Clinical note-taking	6	2	1	1	0	0	0
Generating clinical referrals	3	0	0	0	0	0	0
Enhancing surgical operations	3	3	1	1	0	0	0
Biomedical data mining	2	0	0	0	0	0	0
Generating billing codes	1	0	0	0	0	0	0
Writing prescriptions	1	0	0	0	0	0	0
NLP and NLU tasks							
Question answering	398	194	71	61	54	14	5
Text classification	29	10	6	5	10	2	0
Information extraction	29	12	8	5	4	6	0
Summarization	29	21	7	3	8	0	1
Conversational dialogue	6	6	1	1	5	1	0
Translation	5	1	2	2	1	2	0
Dimension of evaluation							
	Accuracy	Comprehensiveness	Factuality	Robustness	Fairness, bias, and toxicity evaluation	Deployment metrics	Calibration and uncertainty

White space opportunity to build agency



Original Investigation | AI in Medicine

October 15, 2024

Testing and Evaluation of Health Care Applications of Large Language Models A Systematic Review

Suhena Rishi, BA¹; Yuting Liu, MA²; Lucy Qin-Ewing, BA³; et al

> Author Affiliations | Article Information

JAMA. Published online October 15, 2024. doi:10.1001/jama.2024.21700

<https://jamanetwork.com/journals/jama/fullarticle/2825147>


Dimension of evaluation

Thoughtful balance of risk & rewards


50%+ of tasks in a hospital are admin

Share of admin tasks along the treatment pathway¹

(as % of total Full-Time-Equivalents (FTE), sum = FTEs)



Typical admin tasks in hospitals (in %)



Low

Type of Task

Clinician-Facing
Administrative

Clinician-Facing
Clinical

Patient-Facing
Administrative

Patient-Facing
Clinical

Risk Level

High

Essence of Agents


"Agent" can be defined in several ways

- As fully autonomous systems that operate independently over extended periods, using various tools to accomplish complex tasks.
- As workflows that describe more prescriptive implementations where LLMs and tools are orchestrated through predefined paths

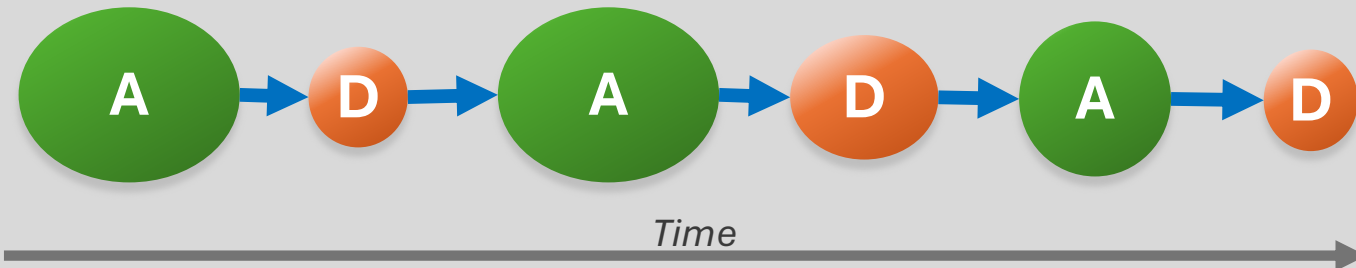
Agents are systems
which provide and preserve agency
to providers

Anatomy of Agency


A = Action, D = Decision/Reasoning



Most pathways
(clinical, operational)
are a sequence of
actions (A) and decisions (D)



Digital tools tend to
increase A and reduce D



Agentic systems should
decrease A and increase D



Time to Think (Tik)
Time to Care (ToC)

What should agents NOT do to undermine Agency?

Don't add WIMPs (Windows, Icons, Menu, Pointers)

- ✓ Agents are not apps
- ✓ No new screens, log-ins, forms and dashboards to navigate
- ✓ Interactions are designed to be natural and considerate, reducing digital friction.

Don't break provider pathways

- ✓ Don't add new actions in existing workflows
- ✓ Work in background mostly and surface when needed integrate into existing

Don't increase the risk posture of care

- ✓ Only automate actions which are automatable
- ✓ Work in sand-boxed, tightly scoped functions

What should agents DO to augment Agency?

Clarity & Consistency


- ✓ Work on domain-limited, real-world problems
- ✓ Provide local intelligence
- ✓ Low resistance, high impact

Transparent Reasoning

- ✓ Feed-back and notifications
- ✓ Work in background mostly and surface when needed integrate into existing

Controlled Autonomy

- ✓ Define intent and reveal action
- ✓ Manage risk (autonomy slider)



*giving provider their most valuable
currency: **Time.***

Time to think. Time to care.

Your Autonomy Slider

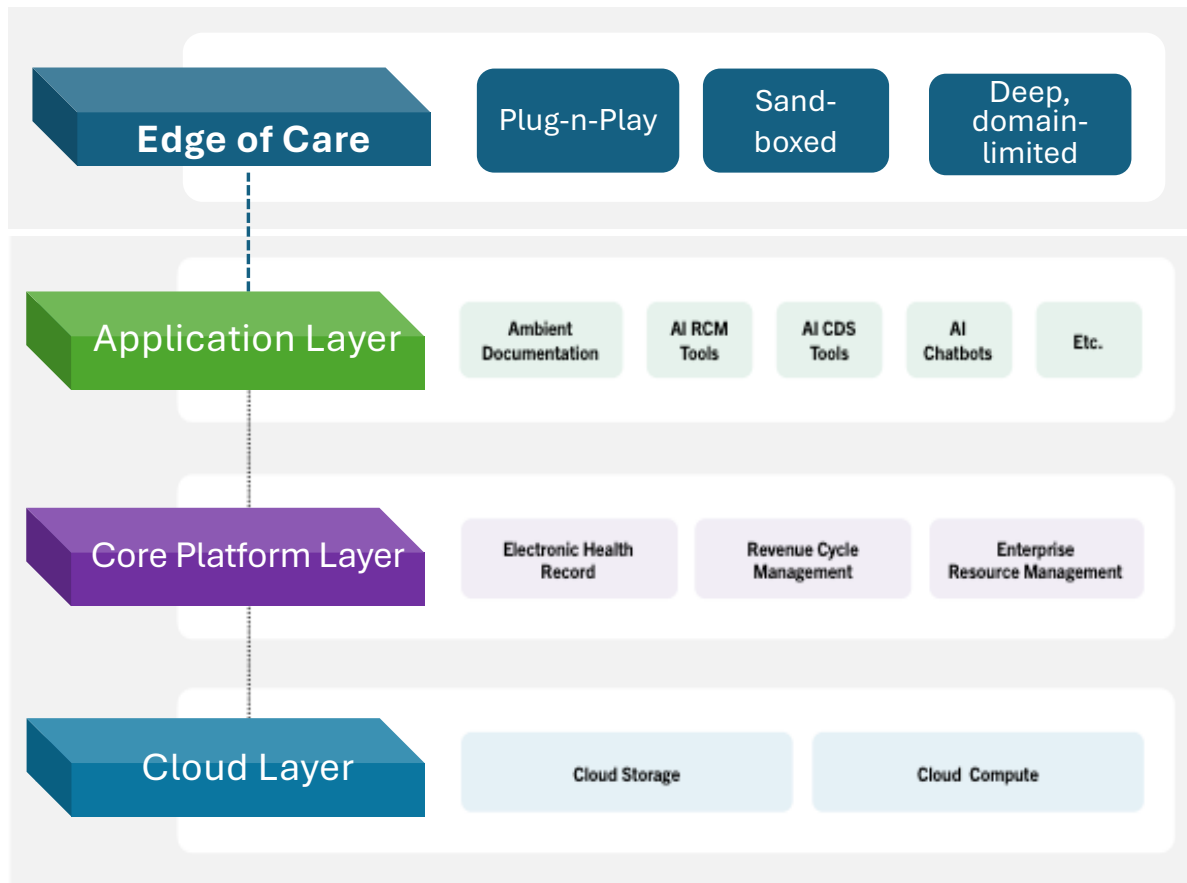
Providers take the controls. Adjust the level of agent autonomy to match your comfort and the task at hand.



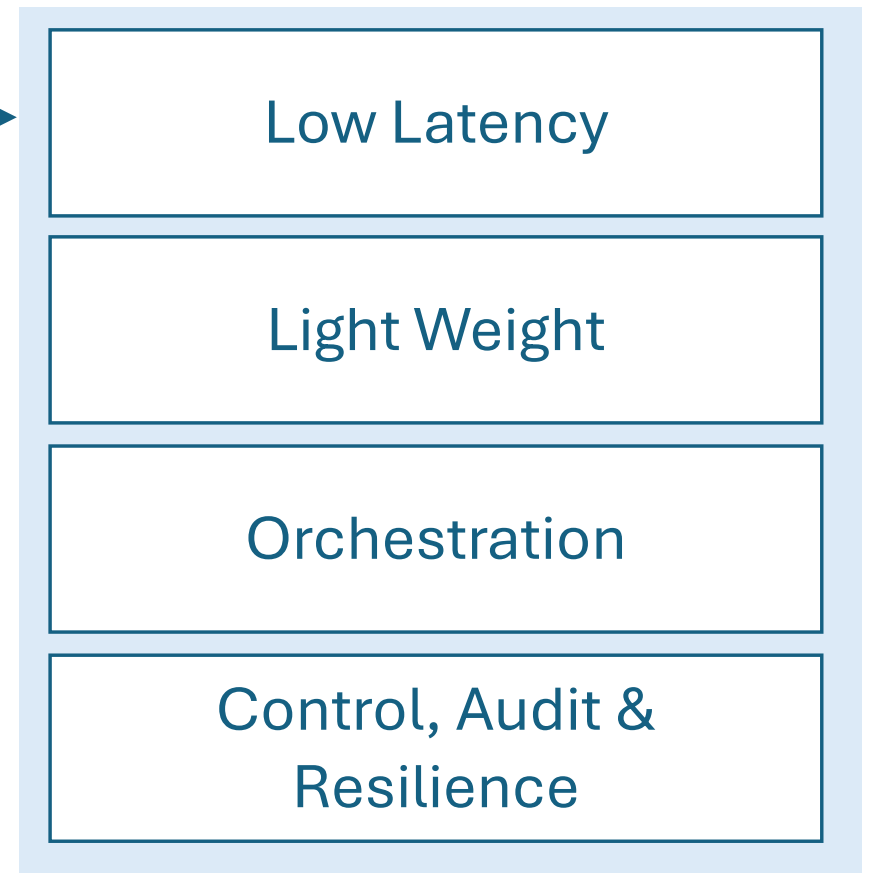
See Disha in Action

How can we preserve agency - edge of care

"Mine the Gaps" at the edge of care



Advantage, Agents



How can we preserve agency - hyper-modal interface


How Natural language performs on the speed dimension

RECEIVE		SEND	
Read	~250wpm	Write	~60wpm
Listen	~270wpm	Speak	~150wpm

Bottleneck

We form thoughts at 1,000-3,000 words per minute

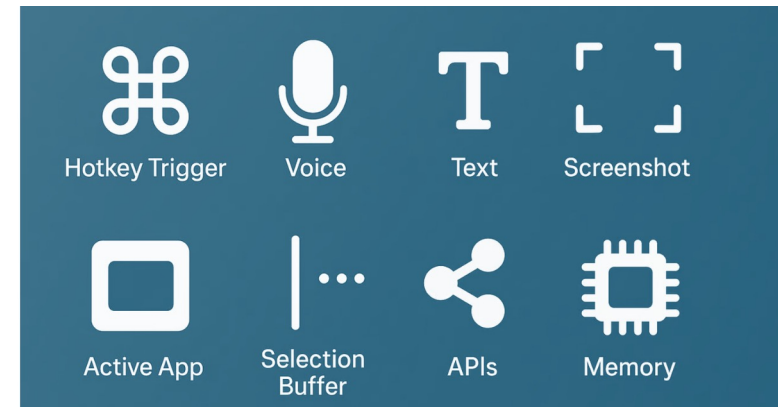
Humans use different modalities for I/O



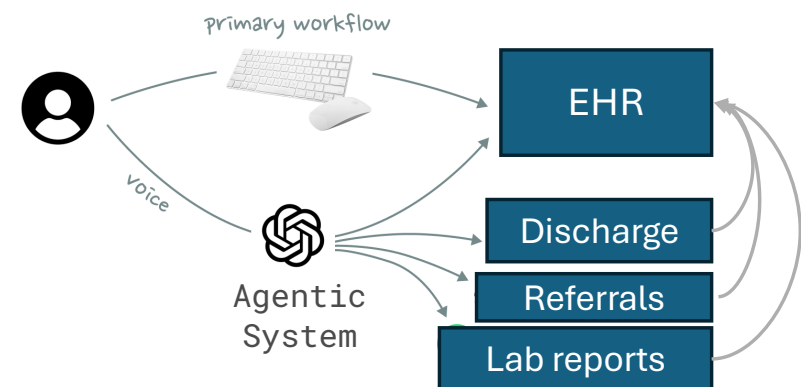
For example, instead of saying "I think what you just said is a great idea", I can just give you a thumbs up. Or nod my head. Or simply smile.

<https://julian.digital/2025/03/27/the-case-against-conversational-interfaces/>
<https://roadtoartificia.com/p/the-hypermodal-interface>

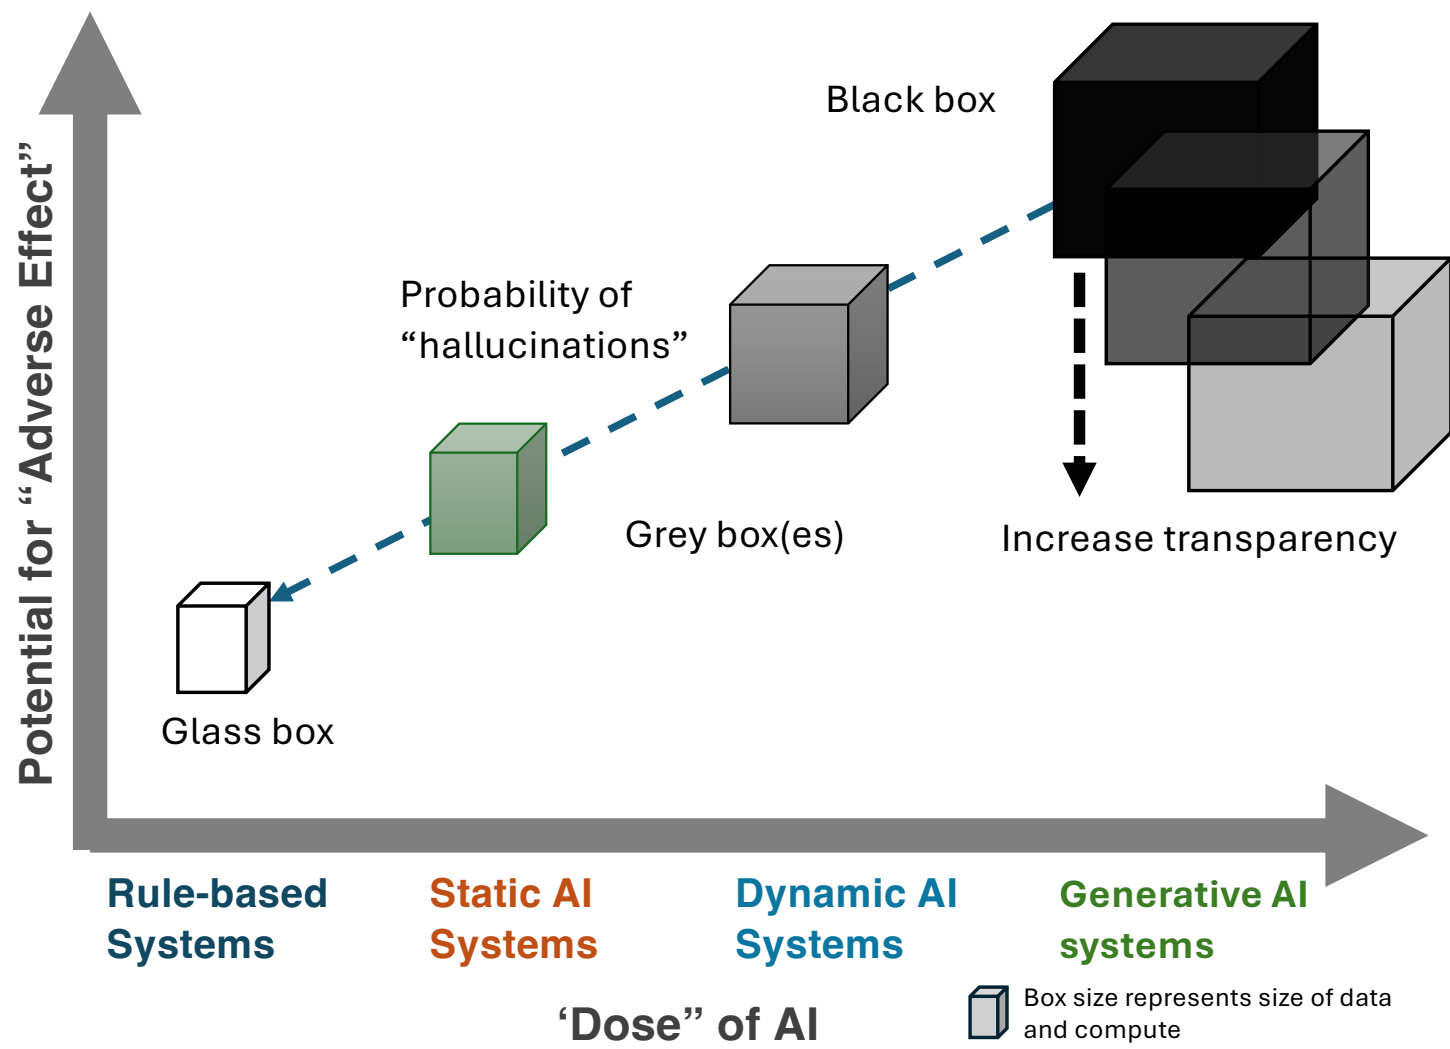
Agents need to interface on all modes



And increase the bandwidth of productivity & collaboration



How can we preserve agency - Err on side of Transparency



Transparent Reasoning

Always know the "Why" behind an action.

Actions are based on reasoning and shaped by your behavior, not hard-coded rules. A "Chain of Thoughts" allows you to build intuition and trust.

R-Tuning: Instructing Large Language Models to Say 'I Don't Know'

Hanning Zhang*, Shizhe Diao*, Yong Lin*, Yi R. Fung*, Qing Lian*, Xingyao Wang*, Yangyi Chen*, Heng Ji*, Tong Zhang*
*The Hong Kong University of Science and Technology
†University of Illinois Urbana-Champaign
(hzhango, sdiao, yilinf, qlianab, tongzhang)@ust.hk
(yifung2, xingyao6, yangyi3, hengji)@illinois.edu

Abstract

Large language models (LLMs) have revolutionized numerous domains with their impressive performance but still face their challenges. A predominant issue is the propensity for these models to generate non-existent facts, a concern termed *hallucination*. Our research is motivated by the observation that previous instruction tuning methods force the model to complete a sentence no matter whether the model knows the knowledge or not. When the question is out of the parametric knowledge, it will try to make up something and fail to indicate when it lacks knowledge. In this paper, we present a new approach called *Refusal-Aware Instruction Tuning (R-Tuning)*. This approach is formalized by first identifying the disparity in knowledge encompassed by pre-trained parameters compared to that of instruction tuning data. Then, we construct the refusal-aware data based on the knowledge intersection, to tune LLMs to refrain from responding to questions beyond its parametric knowledge. Experimental results demonstrate R-Tuning effectively improves a model's ability to answer known questions and refrain from answering unknown questions. Furthermore, when tested on out-of-domain datasets, the refusal ability was found to be a meta-skill that could be generalized to other tasks. Further analysis surprisingly finds that learning the uncertainty results in better calibration and an improved ability to estimate the uncertainty than uncertainty-based testing.¹

1 Introduction


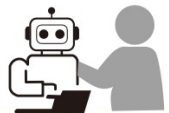



Large language models (LLMs) have demonstrated remarkable performance across numerous tasks; however, they are also plagued by various issues, such as the propensity of large models to fabricate non-existent facts, a phenomenon commonly referred to as *hallucination* (Maynez et al., 2020a).

¹Equal Contribution.
²Our code is available at <https://github.com/shizhediao/R-Tuning>.

Figure 1: An illustration of the parametric knowledge distribution and the instruction tuning data distribution. Pre-training embeds a large volume of parametric knowledge, while fine-tuning may involve knowledge that is not necessarily in the parametric knowledge. We explore the benefits of differentiating instruction tuning data based on parametric knowledge.

Learning, Testing & Evaluation (LTE) is key

Human Agency Scale (HAS) to quantify the team dynamics and degree of human involvement required

					
	HAS H1	HAS H2	HAS H3	HAS H4	HAS H5
Team Dynamics	AI Agent Drives Task Completion The AI agent takes primary responsibility for task execution with no or minimal human oversight.		Equal Partnership The human and the AI agent collaborate closely throughout the task.	Human Drives Task Completion The human takes primary responsibility for task execution with varying levels of AI assistance.	
Required Human Involvement	AI agent handles the task entirely on its own without your involvement.	AI agent needs your input at a few key points to achieve better task performance.	AI agent and you work together to outperform either alone.	AI agent needs your input to successfully complete the task.	Task completion fully relies on your involvement.
AI Role	Automation AI replaces human capabilities		Augmentation AI enhances human capabilities		
Example Tasks	<ul style="list-style-type: none">• Transcribe data to worksheets and enter data into computer.• Run monthly network reports.		<ul style="list-style-type: none">• Create core game features, including storylines, role-play mechanics, etc.• Compile and analyze experimental data and adjust experimental designs as necessary.	<ul style="list-style-type: none">• Coordinate and direct the financial planning, budgeting, procurement, or investment activities.• Design, plan, organize, or direct orientation and training programs.	<ul style="list-style-type: none">• Participate in online forums or conferences to stay abreast of online retailing trends, techniques, or security threats.

Future of Work with AI Agents: Auditing Automation and Augmentation Potential across the U.S. Workforce

Yijia Shao*, Humishka Zope*, Yucheng Jiang, Jiaxin Pei, David Nguyen,
Erik Brynjolfsson, Diyi Yang
Stanford University
{shaoyj, diyi}@stanford.edu

Abstract

The rapid rise of compound AI systems (a.k.a. AI agents) is reshaping the labor market, raising concerns about job displacement, diminished human agency, and overreliance on automation. Yet, we lack a systematic understanding of the evolving landscape. In this paper, we address this gap by introducing a novel auditing framework to assess which occupational tasks workers want AI agents to automate or augment, and how those desires align with the current technological capabilities. Our framework features an audio-enhanced mini-interview to capture nuanced worker desires and introduces the Human Agency Scale (HAS) as a shared language to quantify the preferred level of human involvement. Using this framework, we construct the WORKBank database, building on the U.S. Department of Labor's O*NET database, to capture preferences from 1,500 domain workers and capability assessments from AI experts across over 844 tasks spanning 104 occupations. Jointly considering the desire and technological capability divides tasks in WORKBank into four zones: Automation "Green Light" Zone, Automation "Red Light" Zone, R&D Opportunity Zone, Low Priority Zone. This highlights critical mismatches and opportunities for AI agent development. Moving beyond a simple automate-or-not dichotomy, our results reveal diverse HAS profiles across occupations, reflecting heterogeneous expectations for human involvement. Moreover, our study offers early signals of how AI agent integration may reshape the core human competencies, shifting from information-focused skills to interpersonal ones. These findings underscore the importance of aligning AI agent development with human desires and preparing workers for evolving workplace dynamics.

1 Introduction

Rapid advances in foundation models, such as large language models (LLMs), has catalyzed growing interest in AI agents: goal-directed systems equipped with tool access and multi-step execution capabilities. Unlike standalone models, these agents can perform complex workflows and are increasingly positioned to take on roles across a broad range of professional domains (Jiang et al., 2024; Shao et al., 2024a; Wang et al., 2024b; Yang et al., 2024; Yao et al., 2024). Their integration into occupational settings is already beginning to shape the labor market (Demirci et al., 2025; Hoffmann et al., 2024). For example, research indicates that around 80% of U.S. workers may see LLMs affect at least 10% of their tasks, with 19% potentially seeing over half impacted (Eloundou et al., 2023). Usage data from Anthropic indicates that in early 2025, at least some workers in 36% of occupations already were using AI for at least 25% of their tasks (Handa et al., 2025). While AI adoption in the workplace has shown promise in boosting productivity, it also raises concerns about job displacement, reduced human agency, and overreliance on automation (Hazza et al.,

*Equal Contribution

How can we preserve agency - preserving culture & context

A study* tested major AI Models - the foundations powering tools millions use daily - against cultural values from 107 countries worldwide

English is the reserve currency of AI


Each one reflected the same assumptions - those of English-speaking, Western European societies. None aligned with how people in Africa, Latin America, or the Middle East actually build trust, show respect, or resolve conflicts.

Agentic Systems, especially in healthcare need

cultural intelligence


AI Models and their cultural alignment

This graph visualises the cultural alignment of various LLMs by plotting them on the Inglehart-Welzel Cultural Map - comparing their cultural values with those of different countries & cultural regions.



* https://www.linkedin.com/posts/teybannerman_theres-something-almost-nobody-is-talking-activity-7358405153139367937-4GVf/

The Healthcare Telescope



From Better Eyes to New Windows

Initially, telescopes were valued simply for magnifying what the eye could see—sharper and farther—but they evolved to detect entirely new phenomena like radio waves or gravitational signatures, redefining observation itself

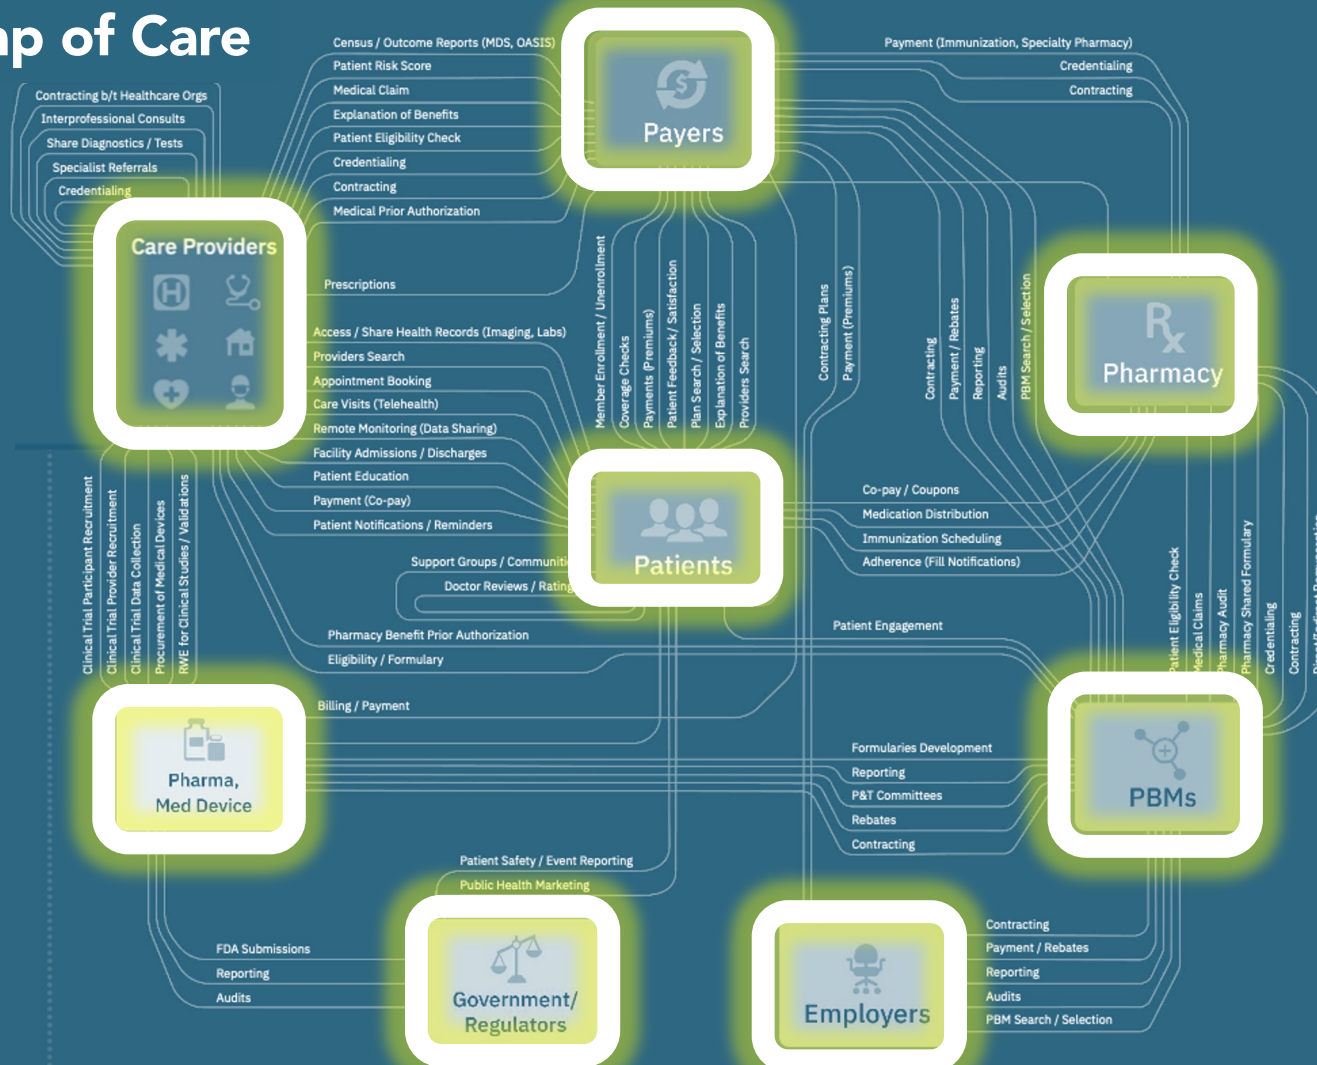
A Shared Direction

The most powerful tools do not just reflect our capabilities back to us. They help us work in new ways, see new patterns, and act with new perspective. To do that, they need to be introduced with care. Not as replacements for human roles, but as extensions of human systems

[The Telescope problem](#)

Increase the aperture of utility in care

The Subway Map of Care





Platform for supply side of health →

Disha (दिशा)
means
"direction"
and
"path"



Happy Providers make Happy Patients

Providers



Patients



in a world
of machines...

Thank You

humanity
becomes the
competitive
advantage.

THANK YOU



**INFOCOMM
MEDIA
DEVELOPMENT
AUTHORITY**