


# Fujitsu Causal AI

Every decision deserves causal rigor



Fujitsu Research of America



## Fujitsu Causal AI

Turning everyday enterprise questions into explainable, defensible decisions in minutes.

# Enterprises Make Thousands Of Decisions; Rigorously Analyze Only A Fraction

**3**<sub>B</sub>

Decisions are made annually by managers and executives collectively<sup>1</sup>.

**38**%

Of C-suite executives would trust AI to make business decisions on their behalf<sup>2</sup>.

**\$250**<sub>M</sub>

Is spent annually by a typical Fortune 500 company on decision making<sup>3</sup>.

## Every Business has a long tail of under-analyzed decisions

- Pricing, promotions, staffing, allocation, and operations decisions happen every week.
- Only a fraction receive deep analysis because expert time is scarce.
- The rest rely on dashboards, correlation, spreadsheets, and gut feel.



**Fujitsu Causal AI**

**Accessible**

**Rigorous**

**Scalable**

1. [Who's making the decisions around here?\(2022\)](#), 2. [SAP Sponsored Survey: Executives Trust AI Over Themselves | SAP News Center \(2025\)](#) 3. [How to make better decisions in the age of urgency | McKinsey & Company \(2023\)](#).

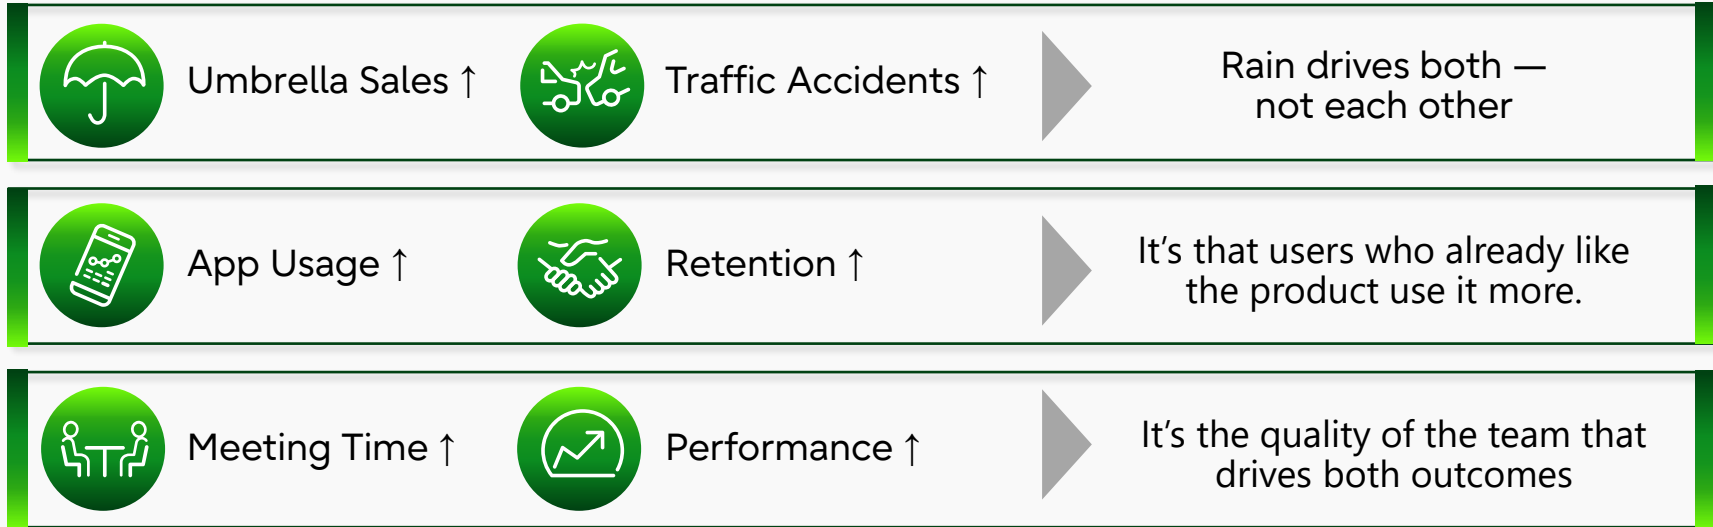
**Would you make a pricing, staffing, or supply-chain decision just because an AI Agent sounded confident?**



- LLMs and agents generate convincing answers, but don't know which lever drives outcomes.
- The challenge is no longer access to AI; but trust in its recommendations.
- Fujitsu Causal AI grounds recommendations in causal reasoning, making them explainable, bounded, and defensible.

# Patterns Tell You What Happened, Causality Tells You What To Change

## Correlation can mislead when the real driver is hidden.



BI dashboards and LLMs detect patterns. Causal inference identifies the lever to pull.

# Fujitsu Causal AI Turns Questions Into Defensible Recommendations



- Prediction is useful, but business leaders need to know what action changes the outcome.
- **Fujitsu Causal AI** is a decision intelligence technology that identifies causes from large-scale data and supports optimal decision-making and actions.


“Causal discovery” and “Causal decision-making” moves teams from insights to action

## 1. Causal discovery

### Key strengths:

1. Fast statistical estimation of causal relationships and causal effects
2. Visualization as a causal graph


A	B	C	...
8.9	7.4	4.2	...
8.4	2.9	9.1	...
0.9	7.3	7.6	...
...	...	...	...



## 2. Causal Decision Making

### Key strengths:

- Recommending actions to move from a current state to a desired state based on causal relationships



# Two Breakthroughs Just Collided

## Agentic AI

- Natural language interaction
- Workflow automation
- Broad accessibility

### Limitations

- Can hallucinate
- Cannot prove causality



## Causal AI

- Identifies true drivers
- Supports intervention analysis
- Produces explainable recommendations

### Limitations

- Requires specialized expertise



## Fujitsu Causal AI


1. Delivers action-ready recommendations **in minutes**
2. Combines insights across **multiple datasets**
3. Identifies true **business drivers**
4. **Recommends actions** under constraints

# Same Analyst, Different Authority

## Before



Business analyst



## After



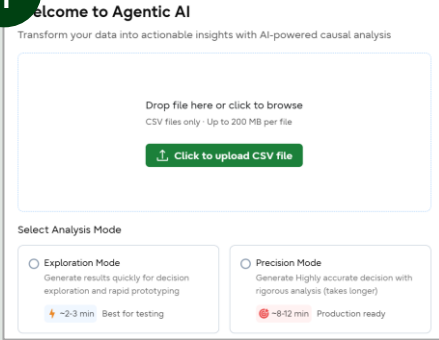
“ I think we should change the price based on this dashboard. ”

“ Increase price by 4% expected margin uplift is 3.1%; constraints are satisfied; here is the causal chain. ”

Fujitsu Causal AI does not replace the analyst.  
It gives the analyst the authority of a senior data science team.

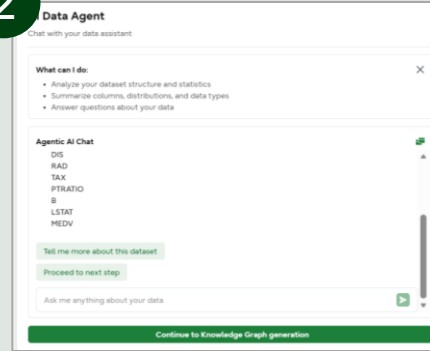
# Demo Walkthrough (Agentic Mode)

1




- Upload your CSV file
- Select Analysis mode

2



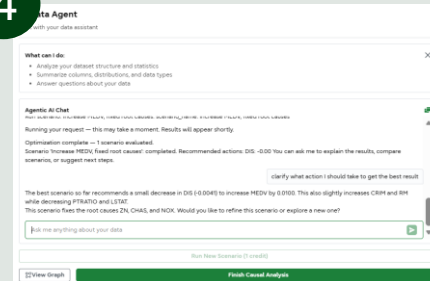
- Ask any questions related to your data set to the AI agent
- Generate a causal graph

3



- Understand the causality
- Start over the graph if you want
- Ask any questions about the results

4



# Demo Walkthrough (Expert Mode)

**1**

Analysis workflow  
Configure, run, and review results.

1. Upload & Preview **In progress** ^  
Upload a CSV or Graph and verify the data was parsed correctly.

2. Discovery Settings **Locked** v  
3. Causal Graph **Locked** v  
4. Optimization Setup **Locked** v  
5. Results & Export **Locked** v

**Discovery algorithm**  
Choose the causal discovery method for this session. This will be used when generating the causal graph from a CSV.

LINGAM  
Single dataset causal discovery.

MultiCa  
Causal discovery across multiple datasets.

TrackR  
Causal discovery guided by a knowledge graph (B2).

**Upload data**  
Upload the required dataset(s) for the selected algorithm. You'll generate a causal graph in later steps.

Drop file here or click to browse  
CSV: Max only: 100x100, 10MB per file  
We'll parse the file and show a preview before running discovery.

[Browse Files](#)

- Select Algorithm and upload your CSV file

**2**

Analysis workflow  
Configure, run, and review results.

1. Upload & Preview **Completed** v  
2. Discovery Settings **In progress** ^  
Configure and run Causal Discovery on your Dataset.

3. Causal Graph **Locked** v  
4. Optimization Setup **Locked** v  
5. Results & Export **Locked** v

**Discovery Settings**  
Configure and run causal discovery on your dataset.

Threshold: 0.01    Gamma: 1    Delta: 2

[Save as New Template](#)    [Run Causal Discovery](#)

- Input "Discovery settings"

**3**

- Understand the causality
- Start over the graph if you want

**4**

**Optimization Setup**  
Set a goal, choose levers, and add constraints to generate a recommendation.

**Output mode**

Explain (single recommendation)  
Returns the best single action.

Enumerate  
Returns multiple candidate actions.

**CRIM**  
0.0063    Target: 0.0063    Goal direction:    Action:    Side effects:    Observed range: 0.063 - 88192

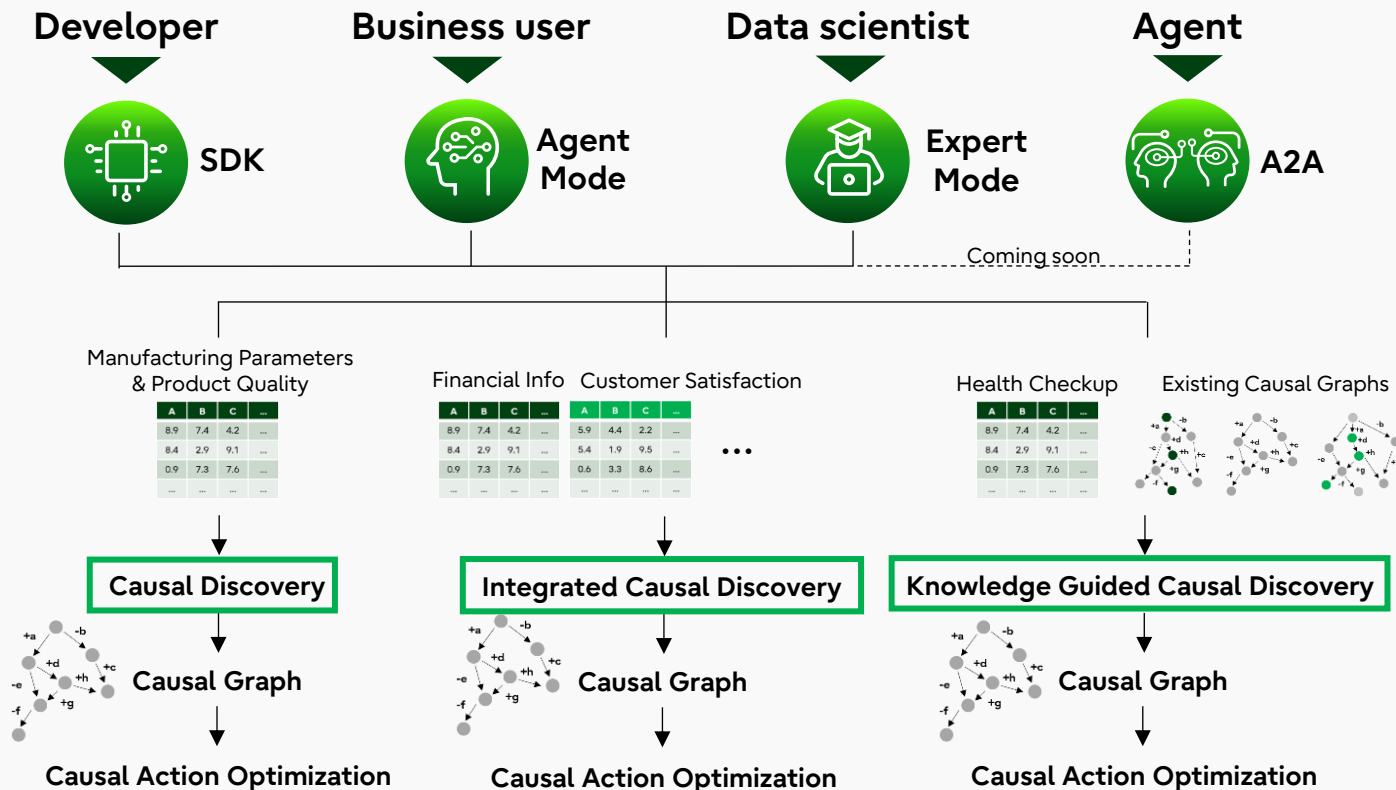
**ZIN**  
18    Target: 18    Goal direction:    Action:    Side effects:    Observed range: 0 - 100

**INDUS**  
2.31    Target: 2.31    Goal direction:    Action:    Side effects:    Observed range: 0.46 - 276

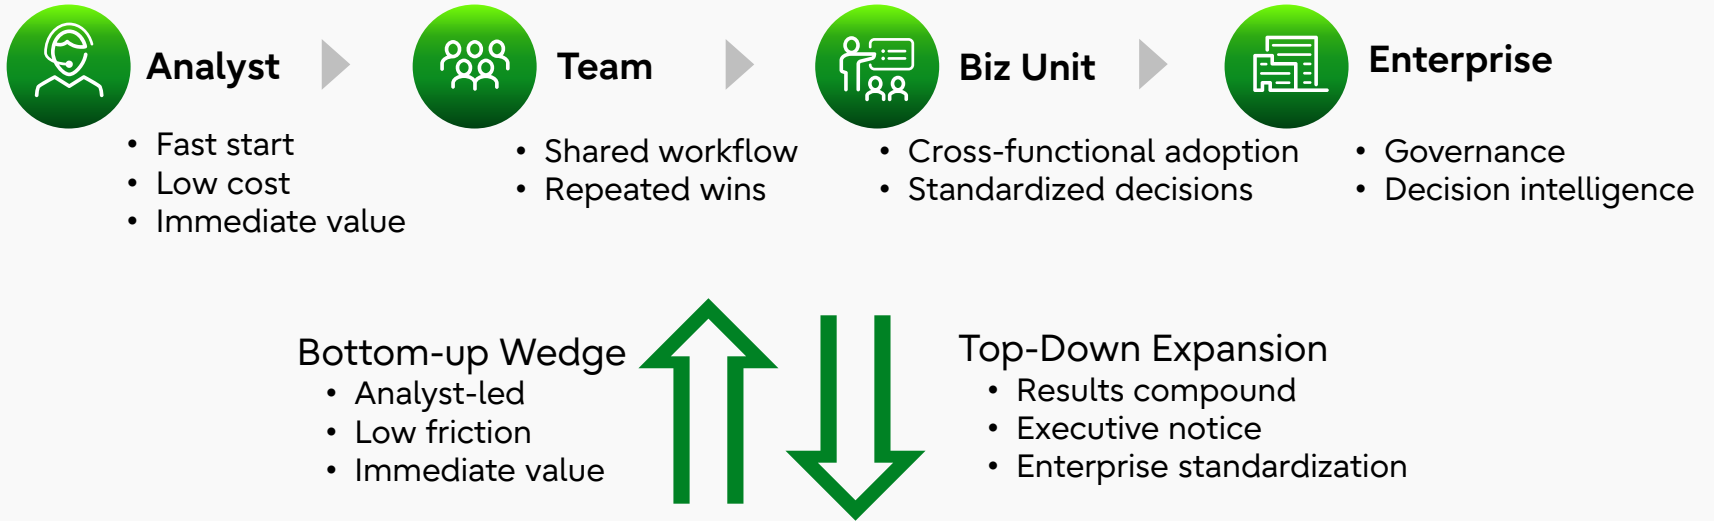
**CHAS**  
0    Target: 0    Goal direction:    Action:    Side effects:    Observed range: 0 - 1

- Optimize setup
- Ensure all inputs are valid
- Get results

# One Causal Engine, Multiple Access Points



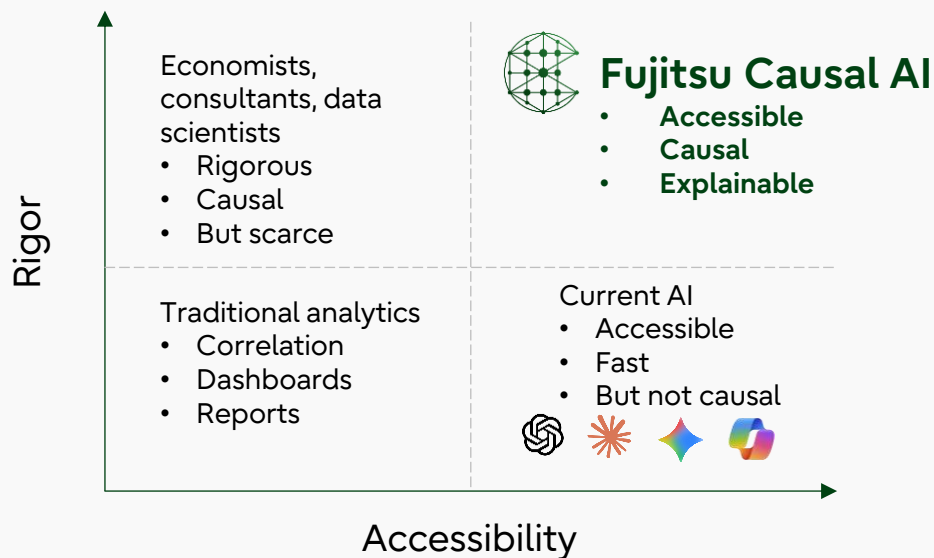
# We Enter Through Analysts, We Expand Through Results



**Every successful decision creates demand for the next one.**

# Causal Rigor Is No Longer Locked Behind Experts

For the first time, organizations no longer have to choose between rigor and accessibility



# Competitive Feature Comparison

- No company offers a feature that automatically generates "action proposals" from scratch
- Few companies implement integrated causal learning and knowledge transfer, and Fujitsu holds an advantage in accuracy
- Causal graph visualization is offered by various companies.

Service	Strategy Proposals	Integrated Causal Learning	Causal Graph Display and Explanation Clarity	Handling of time-series data
Fujitsu	○	○	△	△
causaLens (AI Agents / decisionOS)	△	△	○	○
Microsoft (PyWhy: EconML / DoWhy, Azure ML)	×	×	○ (PyWhy&Azure ML)	×
Google (CausalImpact)	×	△	△	✓
IBM (Causal Inference 360 / Instana)	×	×	✓	△
hootfolio (causal analysis)	△	×	○ (Discretization&missing value imputation)	×
Sony CSL (CALC 4.0) ※ <b>ISID</b> used as consultant	△	×	○	×
xCausal (Woerth) * <b>LTS Consulting</b> , <b>Sigmaxis</b> usage	△	×	○	×
Actable AI	△	×	○	○

# Sample Customer Use Cases

## Distribute optimal coupons

JPN: Major supermarket E (national) and H (regional)

Data partitioning techniques with strong causality




## Brew the ideal beer

swivel & knot

Causal exploration tech for combinations of outcomes

- Aging temperature (°C)
- Fermentation time (min)
- Malting steep water (kg)
- Roasting time (min)
- Input interval (min.)



## Identify financial performance drivers

Fujitsu Corporate Wide Project

Causal discovery: financial and non-financial factors



# Case Studies: Proven Across Industries



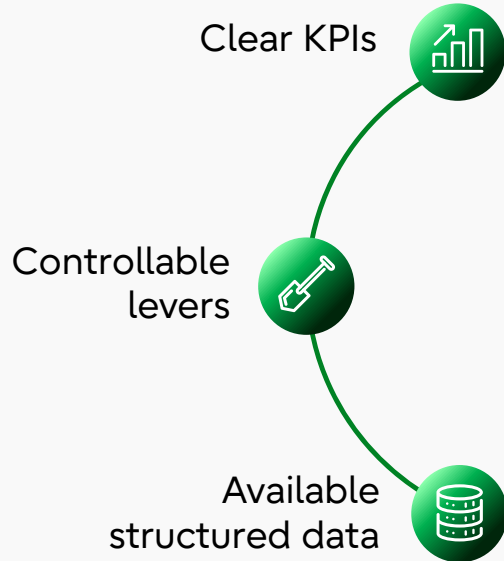
From ESG to healthcare, Fujitsu Causal AI helps organizations uncover hidden drivers, make evidence-based decisions, and accelerate innovation

Collaborator	Domain	Source	Hidden Driver Discovered	Decision Enabled	Key Details
Deloitte Tohmatsu	ESG Management	<a href="#">Press Release</a> , May 29, 2026	Disclosure gaps and competitive differentiation factors	Better ESG strategy	Analyzed ESG disclosure data from 1,000+ listed companies to identify strengths, gaps, and peer differentiation
Kyoto University & Hiroaki University	Population Health	<a href="#">Press Release</a> , March 6, 2025	True causes of sleep disorders	Better health interventions	Combined limited study data with a causal knowledge base built from 20 years of health data and ~3,000 variables
GenQuest	Personalized Healthcare	<a href="#">Press Release</a> , October 9, 2025	Genetic and lifestyle drivers of health outcomes	Personalized recommendations	Revealed how genetic traits, dietary preferences, and lifestyle factors jointly influence health outcomes
Atmonia	Materials Science	<a href="#">Press Release</a> , April 13, 2022	Catalyst characteristics driving performance	Faster materials discovery	Accelerated catalyst exploration using HPC, simulations, and AI to reduce discovery time
Tokyo Medical and Dental University	Drug Discovery	<a href="#">Press Release</a> , March 7, 2022	Causal mechanisms behind drug resistance	Accelerated drug development	Identified previously unknown mechanisms of cancer drug resistance through large-scale causal analysis

# Start With High-Frequency, Measurable Decisions

Focus on decisions you can measure, influence, and repeat

## Best first use cases



## Repeatable decisions



# Adoption Starts With One Decision And Scales With Repeated Value

Start small. Prove value. Scale decision rigor across the enterprise.

**Enterprise:** cross-BU governance, private deployment, expert support.

**Team subscription:** embed into operating rhythm.

**Paid credits:** repeat across weekly team decisions.

**Free trial:** 20 days and 20 decisions.



## Get Started Today!!

1. Free trial (cloud version) now available
  2. Technical consultations available upon request
- Rigor becomes an organizational habit
  - Small wins compound into lasting advantage
  - Better decisions become the default



## Fujitsu Causal AI – Rigor, On Demand

**Thank you**

## High-speed causal discovery using “LayeredLiNGAM” (ECML-PKDD2024)

### Conventional

#### DirectLiNGAM


$x_1$   
↓  
 $x_2$   
↓  
 $x_3$   
↓  
 $x_4$   
↓  
 $x_5$

Bottlenecks in computation when deciding the order of variables

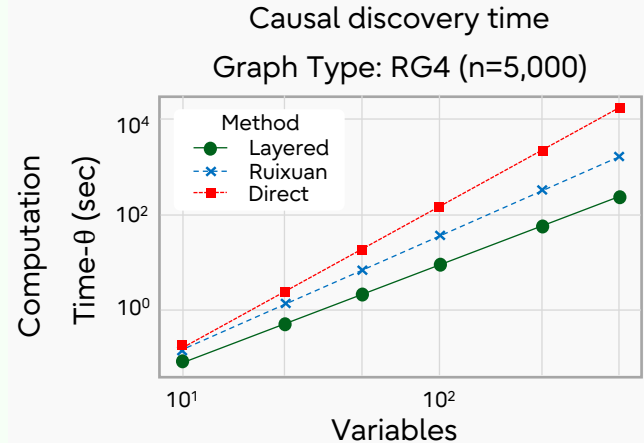


Identify the need for ordering by statistical testing and create “layered groups”

#### LayeredLiNGAM



### Fujitsu Causal AI




## Optimal Causal Decision-making

Recommend the **most effective measures** with the **least negative impact**

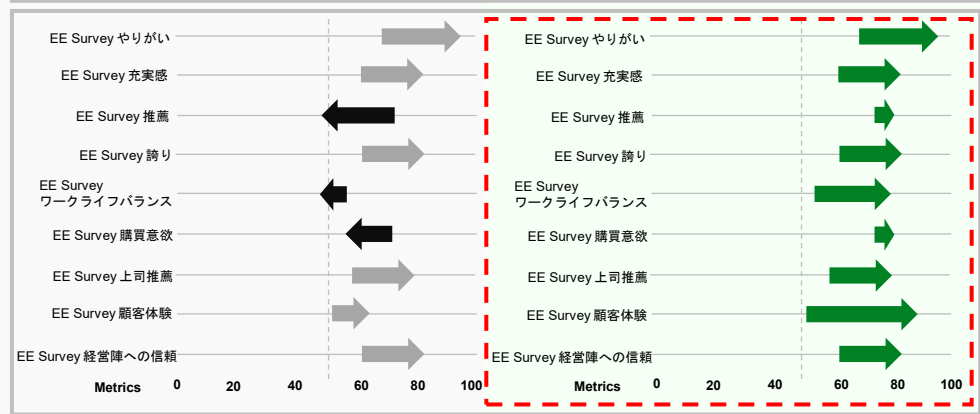
Control the effect of actions on other attributes and optimize by considering multiple causal graphs

- Uses mathematical optimization techniques (**linear programming**) to determine actions that:
  - **Constraints:** Change the value of the target attribute on all causal graphs to the desired value (e.g., no more than 50% turnover risk)
  - **Minimize:** Action Cost and Causal Effect Change in Non-Target Attribute Values

### Conventional



### Fujitsu Causal AI

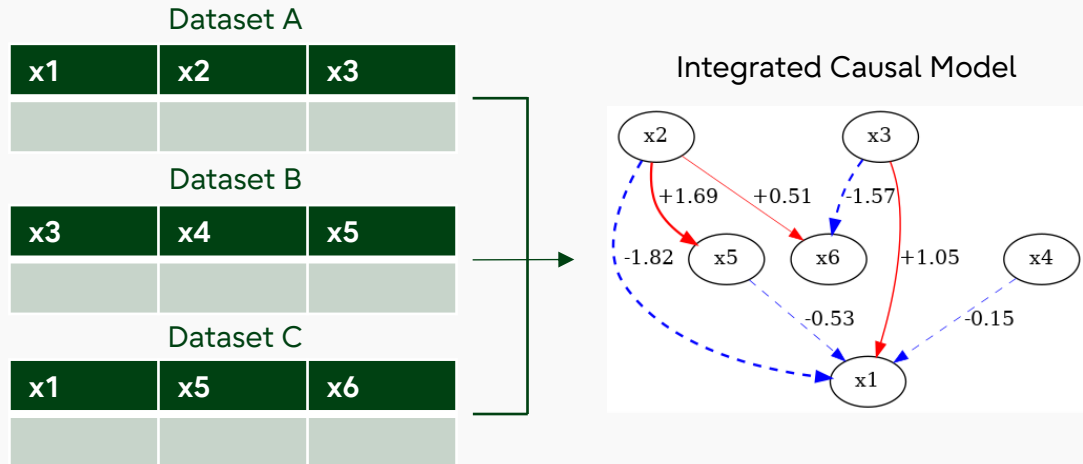


➡ Positive    ⬅ Negative    Achieved targets with no negative impacts

- Purpose: Construct integrated causal models integrating observational data from differing domains
- Reconstructs a single causal relationship from a dataset that has some overlapping variables but also has some distinct variables
  - For example, due to cost issues, each hospital can only measure some of the items
  - Data sets that belong to different domains, such as opportunity data and engagement surveys

Each data set consists of different variables.

Cross-sectional analysis is usually difficult



- Purpose: Use Past Causal Graphs as Prerequisite Knowledge for Statistical Causal Discovery
- Appropriate selection considering differences in **observation conditions** and **observation items**
  - For example, a causal graph estimated from Hirosaki COI health examination data can be used as prerequisite knowledge for causal discovery of customer data to derive appropriate causation.

Hirosaki COI Health Examination Data

Customer Data

Reliable Data

Including data on lifestyle, diet, exercise, and working conditions. (This is hard-to-obtain big data related to health)

