


Fujitsu Causal AI

Every decision deserves causal rigor



Fujitsu



Fujitsu Causal AI

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

Fujitsu was recognized as a Sample Vendor in the **Causal AI category** in **Gartner® "Hype Cycle™ for Artificial Intelligence, 2026"** (Haritha Khandabattu et al., 28 August 2026).
See Gartner disclaimer in appendix.

Enterprises Make Thousands Of Decisions; Rigorously Analyze Only A Fraction

3_B

Decisions are made annually by managers and executives collectively¹.

38%

Of C-suite executives would trust AI to make business decisions on their behalf².

\$250_M

Is spent annually by a typical Fortune 500 company on decision making³.

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

Democratizing

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).

A Confident Agent Is Not The Same As A Trusted Decision

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
- Makes recommendations explainable, bounded, and defensible.

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. ”

With Fujitsu Causal AI, the same analyst now has the authority of a senior data science team behind them.

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

Correlation can mislead when the real driver is hidden.



Umbrella Sales ↑



Traffic Accidents ↑



Rain drives both —
not each other



App Usage ↑



Retention ↑



It's that users who already like
the product use it more.



Meeting Time ↑



Performance ↑



It's the quality of the team that
drives both outcomes

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 Fused



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

Demo Walkthrough (Agentic Mode)

Agentic Mode: Conversational causal analysis for business users

1

Welcome to Agentic AI

Transform your data into actionable insights with AI-powered causal analysis

Drop file here or click to browse
CSV files only · Up to 200 MB per file

Click to upload CSV file

Select Analysis Mode

☐ Exploration Mode

Generate results quickly for decision exploration and rapid prototyping

~2-3 min · Best for testing

☐ Precision Mode

Generate Highly accurate decision with rigorous analysis (takes longer)

~8-12 min · Production ready

- Upload your CSV file
- Select Analysis mode

2

AI Data Agent

Chat with your data assistant

What can I do:

- Analyze your dataset structure and statistics
- Summarize columns, distributions, and data types
- Answer questions about your data

Agentic AI Chat

DISE
RAD
TAX
PFRATIO
B
LSTAT
MEDV

Tell me more about this dataset


Proceed to next step

Ask me anything about your data

Continue to Knowledge Graph generation

- 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

AI Data Agent

Chat with your data assistant

What can I do:

- Analyze your dataset structure and statistics
- Summarize columns, distributions, and data types
- Answer questions about your data

Agentic AI Chat

Run scenario: Increase DISE, reduce RAD, increase TAX, increase PFRATIO, increase B, increase LSTAT, increase MEDV

Running your request — this may take a moment. Results will appear shortly.

Optimization complete — 1 scenario evaluated.

Scenario: Increase MEDV, fixed root causes: completed. Recommended actions: DISE: -0.0001 You can ask me to explain the results, compare scenarios, or suggest next steps.

Clarify what action I should take to get the best result

The best scenario so far recommends a small decrease in DISE (-0.0001) to increase MEDV by 0.0100. This also slightly increases CRIM and BKR while decreasing PFRATIO and LSTAT. This scenario fixes the root causes ZN, CHAS, and NOX. Would you like to refine this scenario or explore a new one?

Ask me anything about your data

Run New Scenario (2 used/5)

View Graph

Finalize Causal Analysis

- Understand and discuss the results
- Modify, improve, and re-run your causal analysis

Demo Walkthrough (Expert Mode)

Expert Mode: Full, fine-grain control for data scientists

1

Analysis workflow
Configure, run, and review results.

1. Upload & Preview **In progress**

2. Discovery Settings **Locked**

3. Causal Graph **Locked**

4. Optimization Setup **Locked**

5. Results & Export **Locked**

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 (KG).

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 files only. Up to 200 MB 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**

2. Discovery Settings **In progress**

3. Causal Graph **Locked**

4. Optimization Setup **Locked**

5. Results & Export **Locked**


Discovery Settings
Configure and run causal discovery on your dataset.

Threshold Gamma Delta

[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.

Variable	Goal direction	Target value	Action	Side effects
CRIM	☐ Minimize ☐ Maximize	0.0063	0.0063	0.0063
ZN	☐ Minimize ☐ Maximize	18	18	18
INDUS	☐ Minimize ☐ Maximize	2.31	2.31	2.31
CHAS	☐ Minimize ☐ Maximize	0	0	0


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

One Causal Engine, Multiple Access Points



Causal Rigor Is No Longer Locked Behind Experts

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



Highlighted Causal AI Case Studies

Brew The Ideal Beer

Swivel & Knot Co., Ltd.



CHALLENGE

Beer quality has traditionally relied on brewers' hard-to-replicate tacit expertise of ingredients and brewing conditions.

APPROACH

Linked brewing parameters (aging, fermentation, malting) to sensory quality, yield & cost via a causal graph.

RESULT

- Delivered a production-ready brewing specification that Swivel & Knot adopted directly for commercial brewing.
- Modeled recipes, processes & sensory data, enabling quantitative trade-offs across quality, yield and cost.

New insights into genetics-lifestyle relationships

Genequest Inc.



CHALLENGE

Understanding how genes and lifestyle habits jointly shape health outcomes is difficult because they rely on complex interactions that defy simple correlation.

APPROACH

Analyzed genetic traits, diet, and lifestyle survey data from ~4,000 people to separate true drivers from mere correlation.

RESULT

- Found that metabolic genetic traits shape beverage preference and physique genes correlate with eating habits and BMI.
- The research will expand to include questionnaire data, health checkup records, and medical information

[Press release: October 9, 2025](#)

Advance Health & Productivity

Fujitsu Corporate-Wide Project



CHALLENGE

Despite treating employee well-being as a strategic priority, Fujitsu had no reliable way to measure how specific initiatives moved business performance.

APPROACH

Applied causal AI to ~100 integrated health, HR & financial data points across roughly 30,000 employees.

RESULT

- Job and life satisfaction were identified as one of the major factors impacting employee leave due to illness
- Launched the commercial "Fujitsu Decision-Making Support Service" built directly on this proof-of-concept.

More 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 Tohmatu	ESG Management	Press Release , 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 & Hirosaki University	Population Health	Press Release , 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
Atmonia	Materials Science	Press Release , 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	Press Release , 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 data consultation and trial: 20 days and 20 decisions.



Get Started Today!

Questions? Email Us!



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

Appendix

Additional Information

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”



Fujitsu Causal AI

LayeredLiNGAM




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



➡ 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



Reliable Data




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




Customer Data



False Hypothesis



Reasonable Hypothesis



Gartner® Hype Cycle™ for Artificial Intelligence, 2026, Haritha Khandabattu et al., 28 August 2026.

GARTNER and HYPE CYCLE are trademarks of Gartner, Inc. and its affiliates.

Gartner does not endorse any company, vendor, product or service depicted in its publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner publications consist of the opinions of Gartner's business and technology insights organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this publication, including any warranties of merchantability or fitness for a particular purpose

Thank you