decisionOS
Data Sheet

causaLens

Overview

decisionOS enables enterprises to pave the way to trusted, collaborative, and actionable decision-making powered by Causal AI. The decisionOS Data Sheet outlines all of the key capabilities of the product:

- Build, deploy, and measure the effectiveness of your decision-making in the real world.
- Identify the cause-and-effect relationships within your data through causal discovery.
- Uncover the mathematical relationships between variables with causal model building.
- Solve complex business challenges with decision intelligence engines.
- Rapidly onboard, collaborate, and publish decision applications, which business stakeholders love.

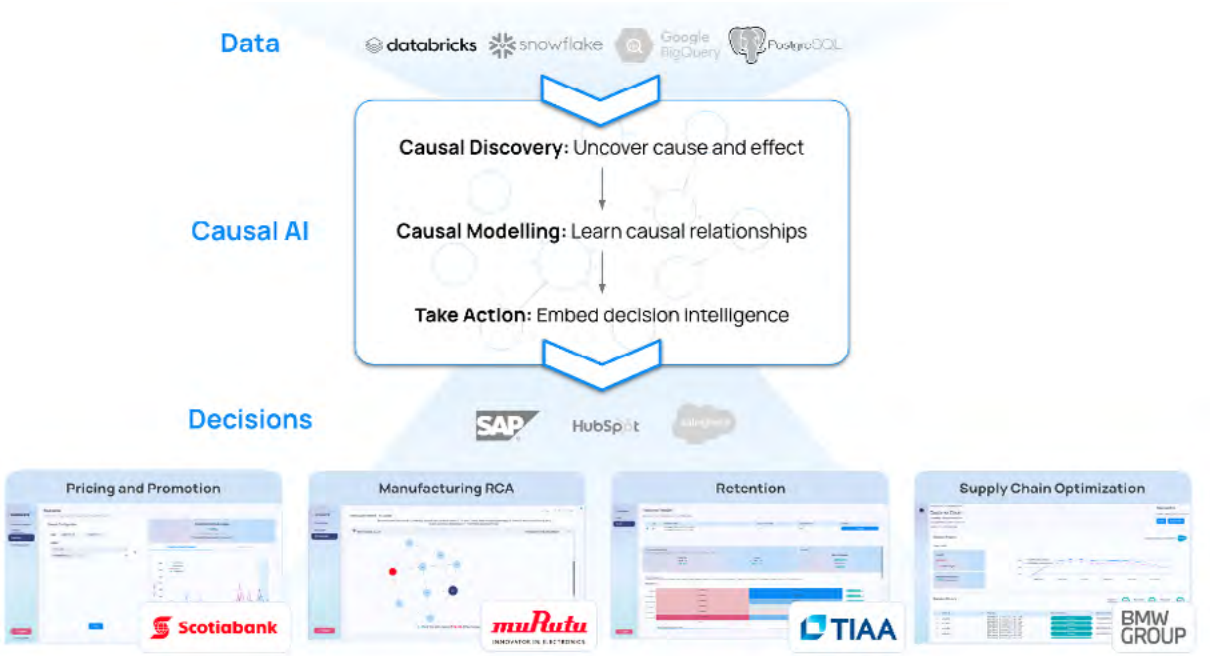

> The causaLens platform has empowered us to create powerful and beautiful decision applications that we rely on for critical decisions."
>
> **Gerald Mullaly, Director Cabinet Office**

> Transparency & explainability of AI models requires an understanding of causality – an inherent advantage of the causaLens platform."
>
> **Wendy Harrington, Chief Data & AI Officer I TIAA**

Causal AI: Why & What

Causal AI is the only technology to reason and make choices as humans do. It utilizes causality to go beyond narrow machine-learning predictions and can be directly integrated into human decision-making. It is the only AI system organizations can trust with their most significant challenges – a revolution in enterprise AI.

Addressing the limitations of traditional ML with Causal AI

Traditional ML:

- Exploits past correlations in data and assumes systems will behave as they did in the past.
- Focuses on predicting outcomes but can't accurately model interventions or counterfactuals.
- Struggles to answer "why" in the form of root cause analysis or (causal) attribution.

With Causal AI:

- Build models that understand cause and effect relationships, going beyond past correlations
- Predict, but more importantly, assess the impact of interventions and counterfactuals, which require a causal understanding of your system.
- Leverage decision intelligence engines to utilize causal models that lead to better decisions.

Beyond traditional causal inference

Traditionally, causal inference has been used in areas where Randomized Control Trials (RCTs) and A/B tests are possible, to infer causal effects. However, recent advancements in the Causal AI field have made it possible to:

- Discover cause-effect relationships from observational data. Read more on the basics of how Casual AI can discover cause-effect here.
- Build structural causal models that can robustly predict and accurately estimate interventions & counterfactuals. Read more on CausalNet, our structural causal model here, and our causal version of decision trees here.
- Discover optimal interventions that maximize KPIs given business constraints and costs using decision intelligence engines like algorithmic recourse.

Improve the explainability of your models & the trust of your business stakeholders

Post hoc explainability methods like SHAP & LIME are typically insufficient.

They can explain what features are associated with the prediction but not necessarily what features drive the outcome. The SHAP documentation warns: "Be careful when interpreting predictive models in search of causal insights."

Causal AI helps Data Scientists get closer to the actual data-generating process and tell a better story for the behavior and mechanics of their model.

Causal AI: Human Guided Causal Discovery

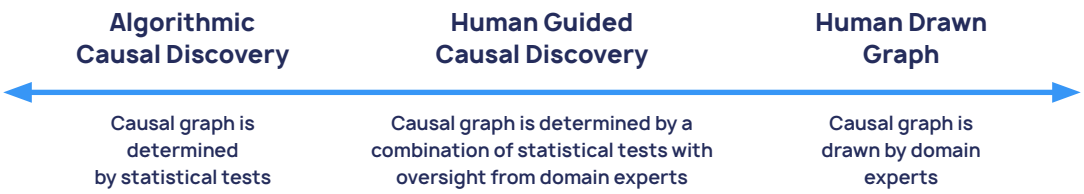

Causal discovery is the process of recovering a causal graph from observed data. decisionOS uniquely supports human-guided causal discovery (HCGD); HGCD combines the best algorithmic causal discovery with your domain expertise, resulting in more accurate causal graphs in less time.

Human-guided causal discovery occurs iteratively, allowing domain expertise to be injected into the discovery process both before and after algorithmic methods have been employed.

This ensures maximal flexibility and allows the discovery process to match your challenges' requirements most naturally.

Algorithmic Causal Discovery

decisionOS supports a range of cutting-edge causal discovery techniques for uncovering causal graphs directly from data.

These methods produce a causal graph, but each algorithm takes a different approach, leading to distinct properties. You can see a summary of the algorithms supported in decisionOS in the table below.

Method	Type	Data Modality Support	Latent Cofounder Identification	Novelty Over Open Source
A*	Score-based	Tabular	No	Yes
A* Global	Score-based	Tabular	No	Yes
A* Local	Score-based	Tabular	No	Yes
FCIComplete	Constraint-based	Tabular	Yes	Yes
FCIExogenous	Constraint-based	Tabular	Yes	Yes
FCITiers	Constraint-based	Tabular	Yes	Yes
Probabilistic FCITiers	Constraint-based	Tabular	Yes	Yes
Parallel-PC	Constraint-based	Tabular	No	Yes
PCMCI+	Constraint-based	Time Series	No	Yes
DYNOTEARS	Continuous Optimization	Time Series & Tabular	No	Yes

Adding Domain Expertise

decisionOS supports the collection of domain expertise through a range of interactive UI components. These provide a convenient interface for non-technical users to add their domain expertise without writing any code. The result is that collecting domain expertise is more collaborative, more accessible, and faster.

Domain expertise can be added using the following interactive UI components:

1. **Tiers:** A hierarchy of features determining which features can causally impact others (e.g., age can cause health problems, but health problems cannot change age).
2. **Edge Encoder:** Add edge constraints to specify whether there is an edge between two nodes (and possibly its direction) or forbid an edge between two nodes. Further, encode edge activations to specify specific types of relationships (e.g., positive linear).
3. **Graph Resolver:** Review existing causal graphs and edit them as required to match the domain understanding of the world (e.g., delete spurious edges).

If a human domain expert is unavailable, or you require quick experimentation, all of these components have variants that are directly integrated into Lexicause, decisionOS' Generative AI toolkit, allowing you to make use of LLM-powered domain expertise.

Data Scientists are not restricted to these components either- there is always the option to add domain expertise directly via Python data structures.

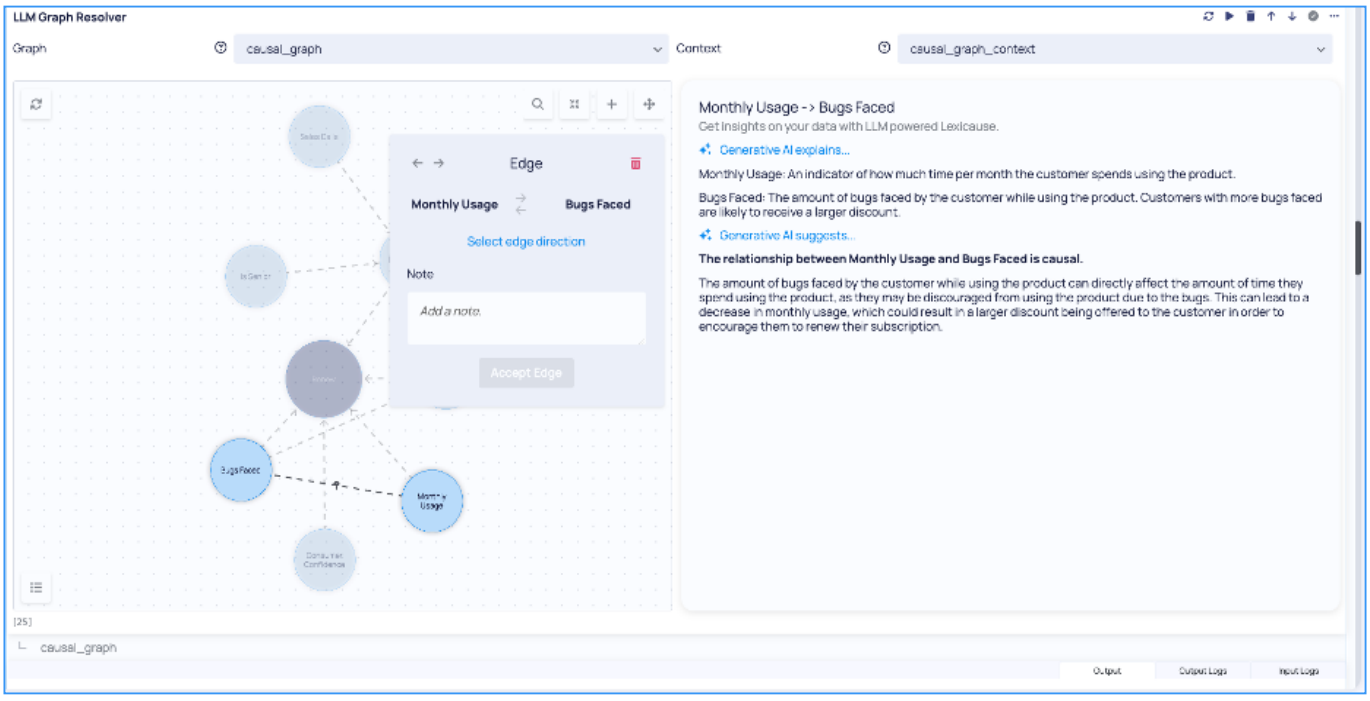

Human-guided causal discovery enhanced with generative AI suggestions.

Causal AI: Causal Model Building

Causal models are unlike traditional correlational machine learning models. Causal models allow predictions to be made about the world given a specific change (interventions) or to predict what would have happened in the past if something else had happened (counterfactuals).

As such, causal models have three key advantages:

- **Advanced "what-if" scenario analysis:** Use interventions and counterfactuals to analyze "what-if" questions accurately and to generate actionable recommendations to drive real outcomes.
- **Robust predictive performance:** Correlational models generalize poorly to unseen data and rely on spurious correlations to make predictions. Meanwhile, causal models exhibit strong out-of-sample performance due to their grounding in invariant causal relationships.
- **Inherent explainability:** Causal models learn the functional relationships between variables based on the cause-and-effect relationships specified in the causal graph. This makes them white box models where humans can inspect and understand every part of the model.

decisionOS comes packed with a library of modeling techniques to deliver value to organizations using causal models as quickly as possible. In addition, decisionOS supports the emerging open-source Causal AI ecosystem by providing interfaces to the most popular libraries and algorithms, for example, DoWhy's own structural causal model.

The causal model inspector component allows you to get an understanding of not only the overall model fit but also the goodness of fit of individual nodes and edges.

1. CausalNet

CausalNet is an advanced structural causal model (SCM) pushing beyond the limitations of traditional machine learning algorithms. CausalNet provides a breadth of functionality, allowing the configuration of models to meet unique needs:

- **Advanced Training Engines:** Select from a range of advanced training engines such as PyTorch, CVXPY, or Pyro to meet performance, data, and compute requirements while smoothing the route to production.
- **Real-World Constraints:** Constrain training to match domain expertise by determining the types of functions that are applied to the edges and nodes.
- **Tabular and Time Series:** Handle tabular and time series data and learn both linear and non-linear relationships.

CausalNet is a powerful and deeply flexible SCM that can be applied to a diverse range of problems effectively.

2. causaLens Decision Trees (CLDTs)

Tree-based models are some of the most popular algorithms today and with good reason. They excel at handling some of the most complex data science cases, easy to understand and interpret, and accommodate both categorical and numerical data.

decisionOS extends tree-based models with the causaLens decision tree (CLDT), allowing delivery of enhanced business benefits:

- **Causal regularization:** CLDTs respect causal relationships, preventing overfitting.
- **Linear leaf models:** The leaves of CLDTs can be fit with linear models, improving performance and enabling explainability.
- **Support for boosting and bagging:** Don't compromise when using CLDTs with full support for advanced methods such as boosting and bagging.

These trees can be integrated into structural causal models using a FullGraphCLDT. This model utilizes decision tree algorithms while providing the capability to do predictive, interventional, and counterfactual modeling.

3. DoubleML

Classical statistical training routines can fail when training Structural Causal Models (SCMs), such as CausalNet, because rather than learning true causal relationships, these often simply learn statistical associations.

To ensure that decisionOS causal models learn the true and unbiased causal effects, decisionOS implements a specialized training routine for SCMs called DoubleML.

4. Open Source

decisionOS provides a highly extensible ecosystem of packages and components. Open-source packages are already deeply integrated across a range of different tasks:

- Discovery of classical machine learning models using libraries such as scikit-learn, XGBoost, LightGBM, PyTorch, and Tensorflow.
- Effect estimation enhanced with EconML and DoWhy.
- Optimization of outcomes using SciPy.

To make it simple to plug and play additional open-source libraries and methods, decisionOS also exposes a standardized interface. This interface is common across models, causal graphs, and decision intelligence engines; therefore, if an external algorithm is integrated within the interface, it can then be utilized as if it were a first-class citizen.

Causal AI: Decision Intelligence Engines

decisionOS provides decision intelligence engines (DIEs) that automate the orchestration of causal models to solve common business challenges. decisionOS has five available DIEs:

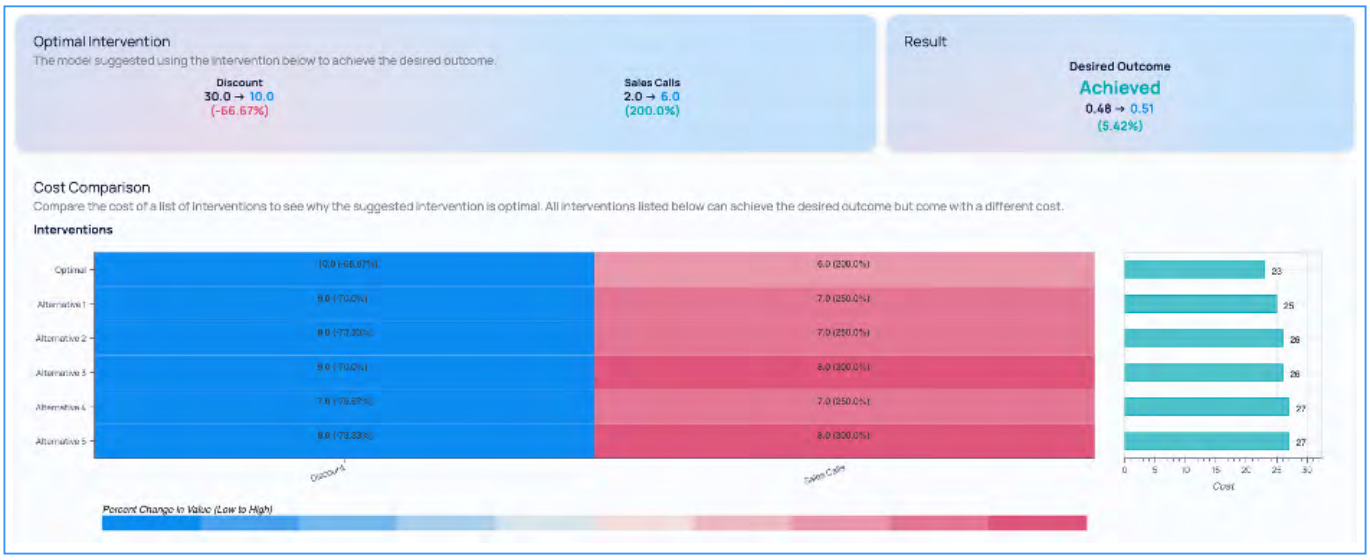

An example showing the different alternative interventions considered during optimization with the lowest cost option being selected.

1. Optimization and Recourse

The optimization and recourse DIEs allow questions such as the following to be answered:

- What action should I take to achieve a specific outcome?
- What action should I take to maximize (or minimize) a specific outcome?

These engines can be configured to respect real-world business with costs and constraints. They help to analyze the downstream effects of actions while evaluating their costs, providing clear suggestions and interpretable justifications for optimal actions.

2. Root Cause Analysis

The root cause analysis (RCA) DIE enables the identification of the true root cause drivers of outcomes. The RCA DIE provides capabilities only found in decisionOS:

- Unique causal approaches ensure that true root causes are identified so attention can be focused on diagnosis without distraction from spurious correlations.
- Causal graphs allow visual representation, enabling complex real-world challenges with large numbers of dependencies to be tackled.
- Built to scale to meet the demands of enterprise workloads from factory floors to global supply chains, handling graphs with hundreds of nodes and thousands of edges.

3. Causal Impact

If a decision has already been made, the causal impact DIE enables the evaluation of the impact of that decision. Causal impact calculates the counterfactual question asking, "What if I had not made the decision?" and displays the change in the relevant KPI or metric given that the decision was, in fact, made. This means that decisionOS can assist in selecting the optimal actions and evaluating the impact of those actions. Causal impact supports the ability to evaluate decisions using a causal model, as well as in a model-free way.

4. Causal Effects

Performing "what-if" analysis using correlation-based machine learning models can be misleading as these models cannot quantify how variable changes influence the desired outcome. Causal effect estimation solves these challenges by providing a set of tools that allow the determination of targeted and specific "what-if" effect estimates as variables are treated with certain actions.

decisionOS' causal effects DIE comes packed with a range of features to support enterprise needs:

- Supports calculating the average, conditional average, heterogenous, or individual treatment effect estimates.
- Brings in meta-learner capabilities that leverage existing correlation-based models within a causal framework to estimate treatment effects accurately.
- If there is existing usage of the open source causal packages, these can plug and play directly with the decisionOS ecosystem.

5. Causal Fairness

As algorithmic decision-making becomes increasingly prevalent, it is critical that model inputs and outputs are carefully monitored to prevent discrimination and de-risk decision-making.

decisonOS provides a range of cutting-edge capabilities to mitigate bias:

- **Fairness requires causality:** Correlation-based fairness metrics fail to capture real-world nuances correctly.
- **Fairness decision algorithm:** Selecting the correct fairness metric can be overwhelming, with conflicting options, which vary from use case to use case. decisionOS provides a simple fairness decision algorithm to guide your selection.
- **Fairness engine:** Ingest a trained causal or non-causal model and generate a suite of fairness metrics to assess bias rapidly.

Example showing the root causes of net financial assets.

An example of the causal impact component demonstrating the estimated effect of Proposition 99, anti-smoking legislation increasing the price of cigarettes, on sales in the state of California.

Decision App Development (Dara)

decisionOS allows Data Scientists to build production-ready applications to power decision-making using Causal AI in pure Python. Expressive interfaces can be created from minimal code, while application development is batteries included: there is no configuration required, and a library of existing applications and components to select from accelerates the time to value.

The application development framework has the following advantages:

- **Extensible:** The application framework has extensibility at its core, and extending the system should be as simple as building an application. Meaning that the existing components can easily be added to with custom components.
- **Smooth learning curve and maintainability:** Going from simple dashboards with no user interaction to complex multi-page apps with lots of interaction points is straightforward while keeping the code maintainable.
- **Native web app performance:** The framework has close to native web app performance and only makes calls back to the Python backend when it really needs to.
- **API access:** The framework supports the creation of simple APIs so that automated processes or advanced users can access the same functionality as the dashboard without multiple solutions.

The application framework provides out-of-the-box components for embedding Causal AI-powered decision intelligence capable of scenario planning, recourse, fairness, effect estimation, and root cause analysis. Empower application users to explore and generate recommendations while building trust through interpretability and explainability.

Native Notebook Integration

Use decisionOS to develop interactive components directly within the notebook environment and render them inline, allowing for advanced and tailored exploration of phenomena seamlessly. Different skill sets within organizations can lead to friction and communication barriers. decisionOS helps to overcome these challenges by providing native APIs to push data and context between notebooks and apps- enabling rapid collaboration between disparate groups.

decisionOS comes packed with a pre-built library of these interactive components called MetaCells. MetaCells have been developed to assist in tasks across the causal workflow, from causal discovery model building to model serving. MetaCells can be added to the notebook environment with the click of a button and then readily configured using simple user interfaces.

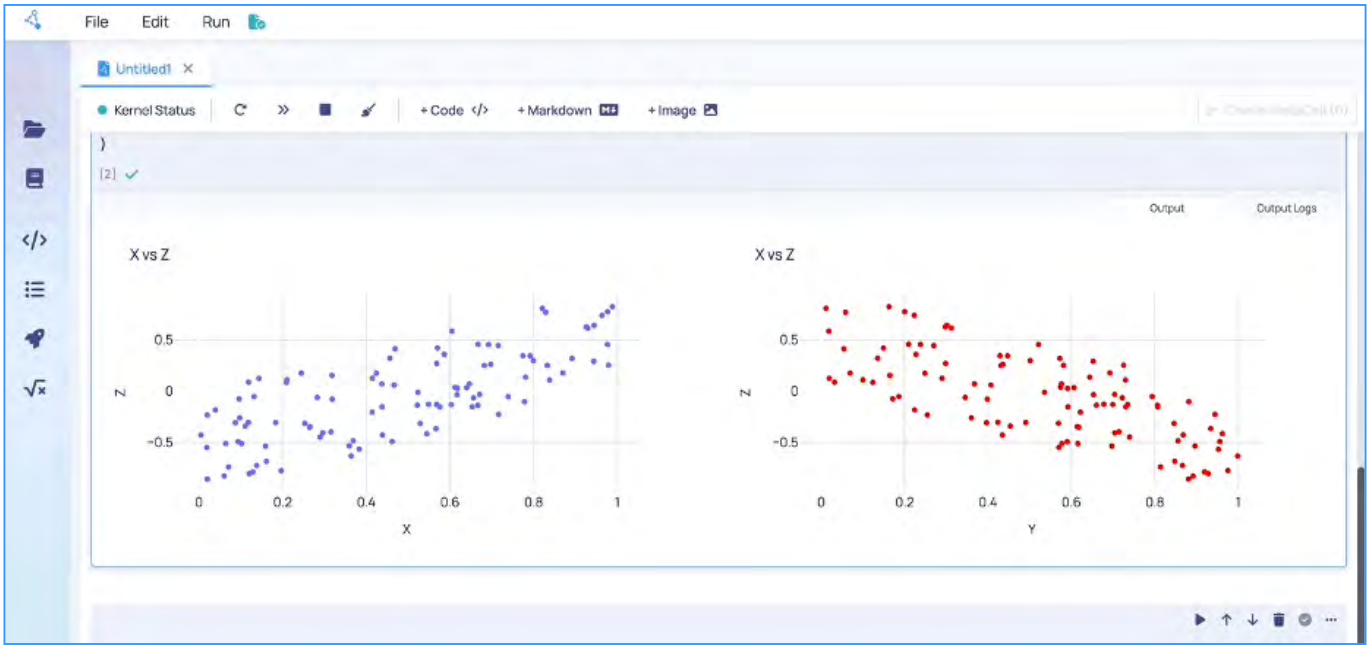

Demonstration of how to add a MetaCell to your notebook.

Application Publishing

Publishing of components and pages from notebooks as standalone applications is as simple as a single button click.

This allows Data Scientists to rapidly move from experimentation workflows into shareable, consumable, stakeholder-facing applications.

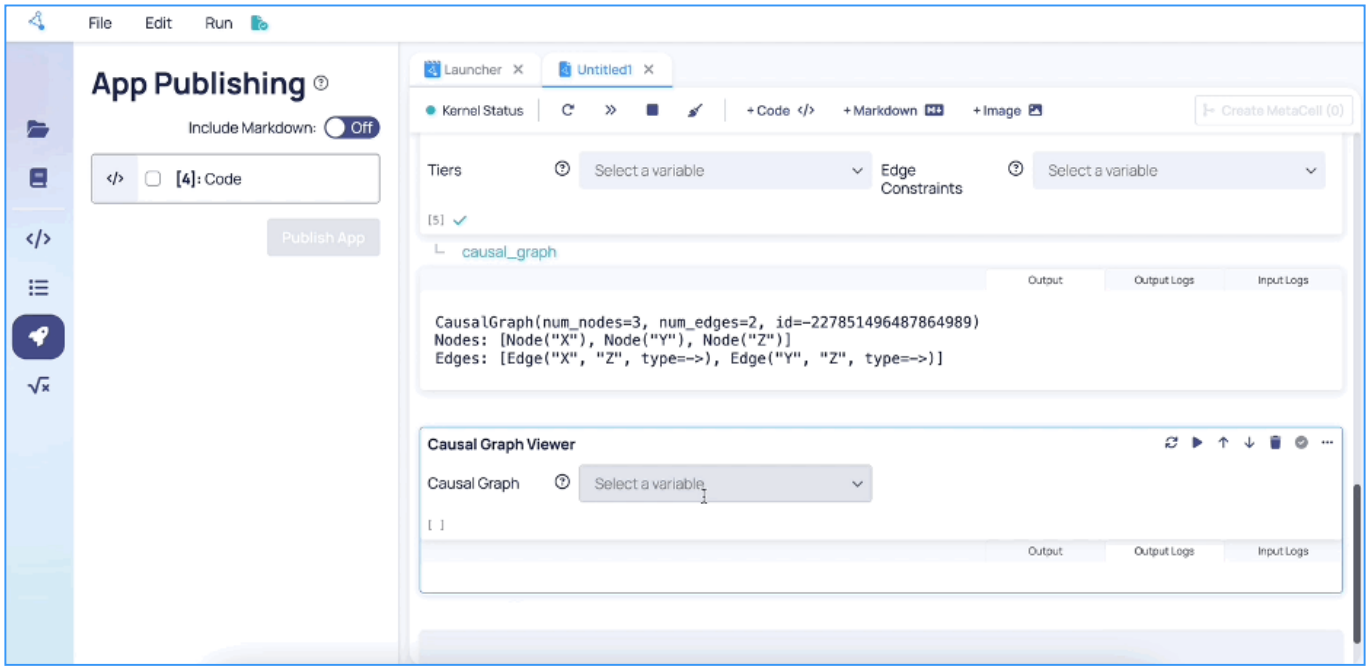

Demonstration showing how to publish interactive outputs as a standalone application.

decisionOps

Use decisionOps to orchestrate workflows, trust the decision-making process, and measure the true causal impact of decision-making.

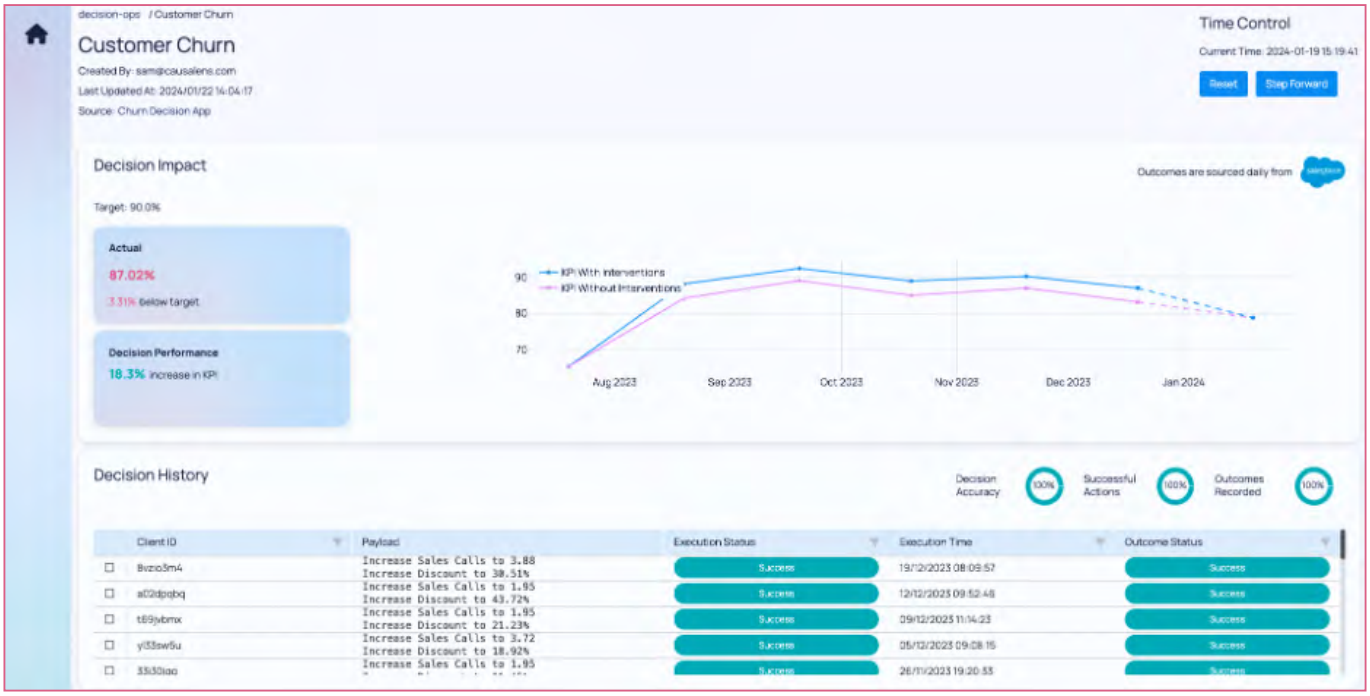

Demonstration showing how to publish interactive outputs as a standalone application.

decisionOps is focused on closing the loop of enterprise decision-making processes in three stages:

1. Recommendation

Causal AI doesn't just generate insights; it provides actionable recommendations. Actionable recommendations allow the implementation of strategies and policies in the real world. decisionOS makes it simple to harness Causal AI to generate actionable recommendations from causal models or decision intelligence engines.

2. Action

decisionOps provides the tools to translate recommendations into the real world. Flexibly and efficiently realize recommendations in systems of record such as ERP or CRM systems with integrations to the most popular technologies- Salesforce, SAP, Oracle, and many more.

3. Traction

In order to understand the impact of actions over time, decisionOps tracks them from implementation all the way to outcome realization. For example, if a retention strategy has been implemented for a segment of customers, it is important to track its progress over time to understand whether it has been successful. decisionOps automates this process by regularly checking updates in the system of record.

Using Causal AI decisionOps uniquely allows inspection of the counterfactual impact of actions: "What would have happened if I had not taken such an action?". This allows decisionOps to go beyond reporting model accuracy. decisionOps can report on decision-making accuracy- allowing the quantification of the impact of the actions being taken in the real world.

Action MetaCell showing how integration with Salesforce can be rapidly configured.

Schematic of the different steps within decisionOps and how it can be used to influence a top-level business KPI.

Causal AI + Gen AI, better together

Generative AI provides a new way for users to interact with applications seamlessly with natural language. Nevertheless, they exhibit certain limitations when employed independently for enterprise decision-making due to issues related to consistency, transparency, and a deeper comprehension of causal relationships within data.

Causal AI, by contrast, inherently offers a transparent framework for revealing these cause-and-effect relationships, building causal models, and offering actionable recommendations that decision-makers can trust.

Together, Generative AI and Causal AI complement one another, helping to ground responses in cause-and-effect reality while maintaining the fluid natural language interface that users are growing accustomed to. decisionOS has three key features that enable this pairing:

Build Better Causal Graphs

causaLens has pioneered human-guided causal discovery that blends domain expertise with the best algorithmic approaches in order to discover cause-effect relationships and causal graphs.

Generative AI enhances this solution by introducing a domain-knowledge assistant to the causal discovery process. The generative model offers directional causality suggestions along with accompanying explanations. This provides Data Scientists with an initial starting point, which enhances productivity.

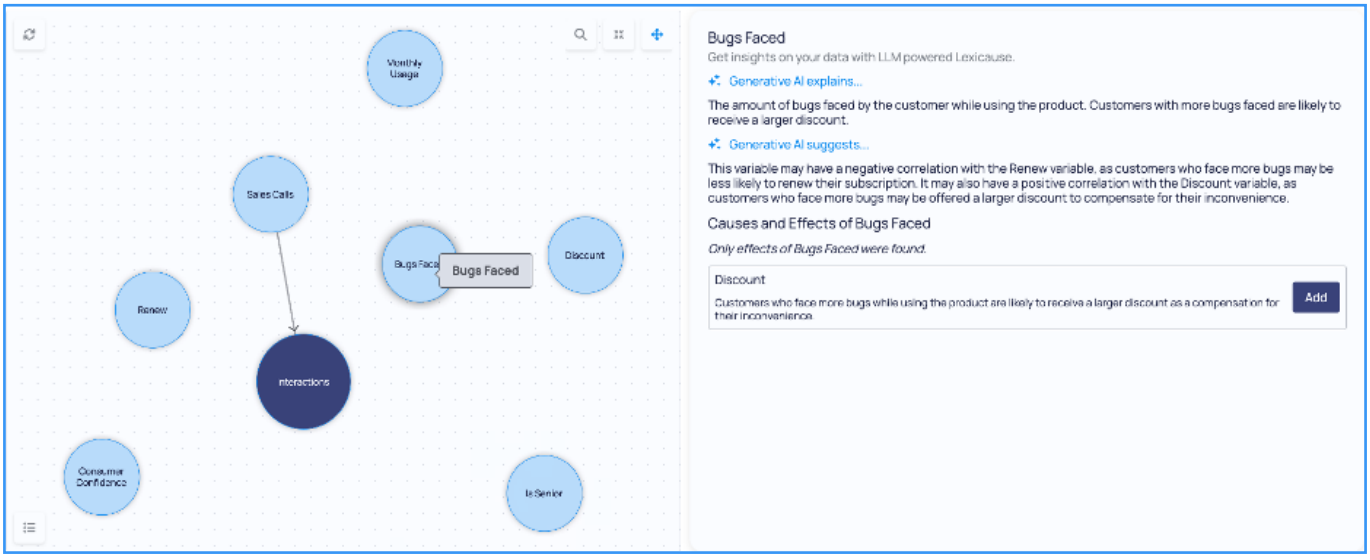

Selecting a node in the Generative AI-powered causal discovery process generates suggestions for causal relationships without the need for a domain expert.

Explain Causal Relationships

Causal graphs are inherently transparent, clearly showing the relationships between variables. Meanwhile, causal models embed the mathematical relationships that flow through the graph.

Through the application of a generative AI layer on top of the causal assets, a generative model can simplify the interpretation of the causal model in a language more accessible to business users. This enhances the accessibility and understanding of these models- gaining trust within the enterprise.

Embed Causal Thinking into LLMs

The integration of human-refined causal assets with large language models provides a synergy and creates responses distinct from those from a generative model in isolation. Causal graphs, causal models, and decision intelligence engine outputs can be passed to the language model, grounding it in reality. This offers the potential to enhance the quality of responses and address some of the trust-related challenges commonly associated with generative models.

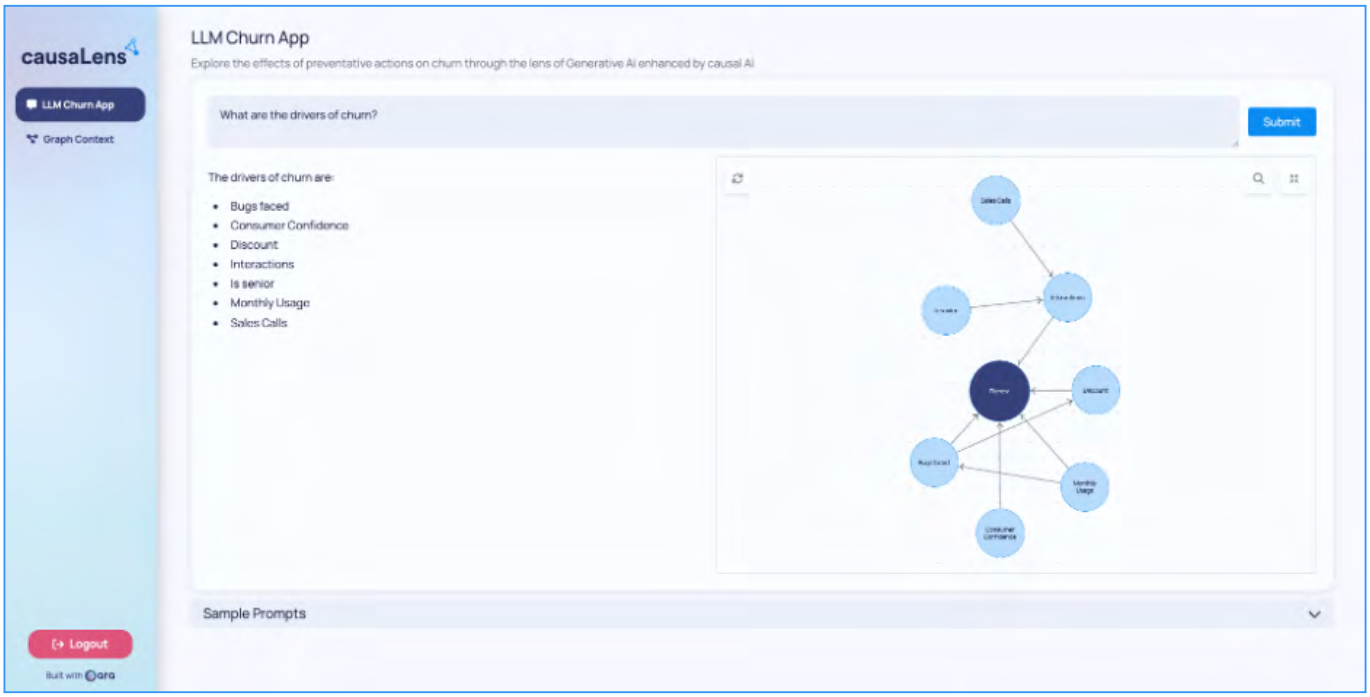

For example, in the analysis of churn-contributing factors, the responses now contain causal drivers, as a representation of causal assets has been embedded within the LLM.

Build vs Buy

Why spend money on a vendor solution when you already have a high-performing Data Science team who could build this themselves?

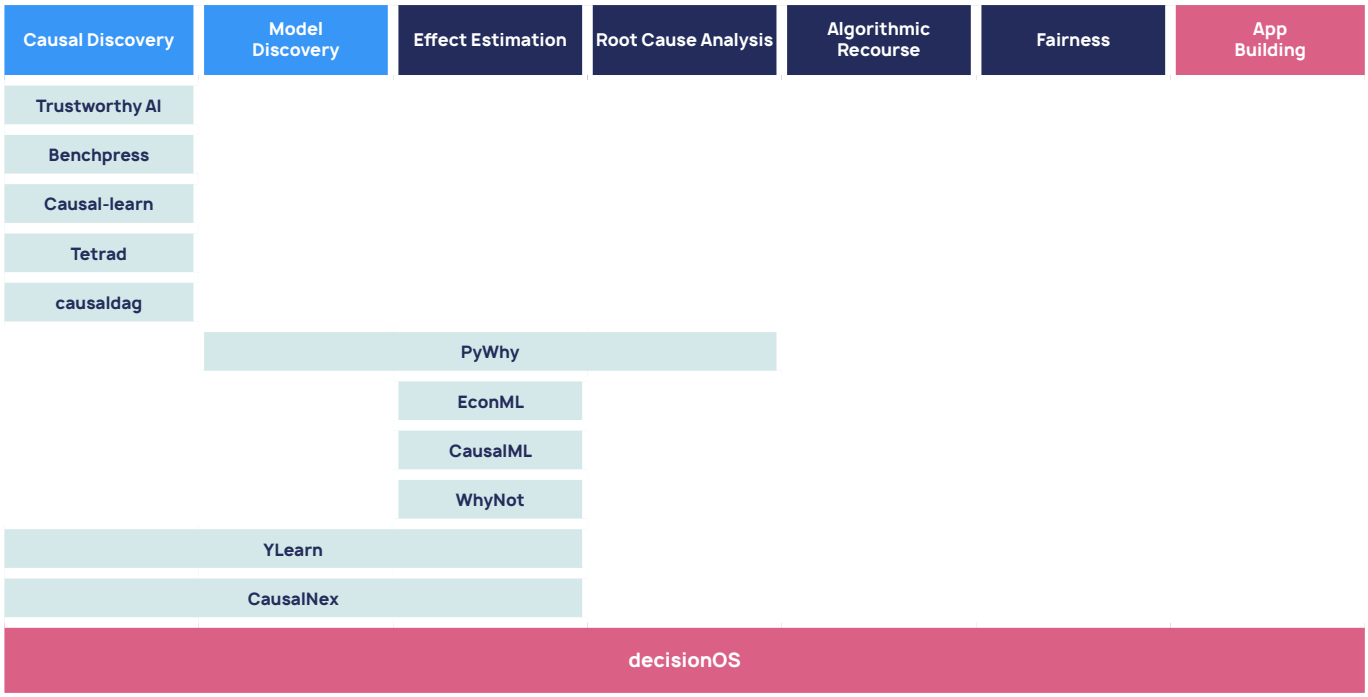

decisionOS covers the entire open-source ecosystem and much more.

An end-to-end solution for going from data to decisions:

decisionOS simplifies and enhances the highly fragmented Causal AI open-source ecosystem, significantly increasing the time to value

Integrate the best of open source with causaLens latest research

decisionOS integrates the best of open source with causaLens' innovations to provide an enterprise-ready Causal AI experience. Building up and maintaining an internal capability is a time-consuming and costly endeavor that detracts from your core company mission.

Leverage our extensive Causal AI research applied to your key use cases

The causaLens team offers over 50 years of experience applying Causal AI to a range of Enterprise use cases. We are familiar with the process of driving value for Enterprises across a range of sectors and have reflected this in many of the components within decisionOS

Extend decisionOS as needed

decisionOS can easily be extended due to the large array of defined interfaces and API-first architecture. It is compatible with any Python library or package

MetaCells + templates

Onboarding and learning a new technology can be difficult. Fortunately, decisionOS comes packed with a range of examples to help get Data Scientists up to speed with Causal AI.

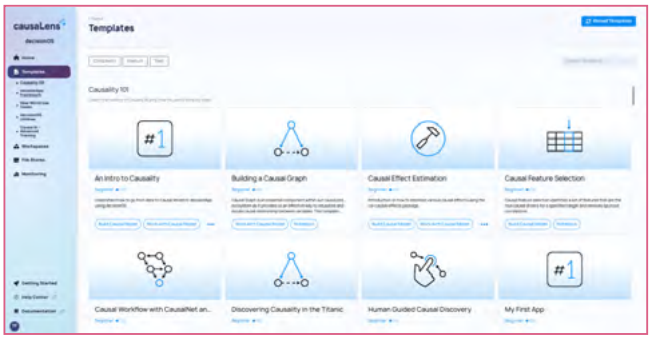

Example showing the root causes of net financial assets.

Templates

Templates are pre-built examples that consist of the following:

- Tutorials educating Data Scientists on how to get started with the causal workflow and application development.
- Examples showcasing how to apply decisionOS technology to real-world datasets end to end.
- Deep dives into more advanced topics and content.

decisionOS has over 50 templates out of the box, and each is launched into an editable environment, allowing them to be tweaked to custom data, use cases, and needs.

Templates accelerate the time to value when using decisionOS by providing a wide range of existing materials as a starting point.

MetaCells

In addition to templates, decisionOS has MetaCells. MetaCells are simply multiple notebook cells merged together. By combining code cells that define visual user interfaces with processing logic, MetaCells provides an easy-to-use user interface for performing common workflow tasks.

decisionOS comes packed with a library of existing MetaCells. These can be dropped into workflows with the click of a button- making it straightforward to get up and running.

MetaCells can be created by any user by simply highlighting two or more code cells and merging them together. Furthermore, any MetaCell can be double-clicked to unwrap it, allowing the underlying code to be edited directly before running the cell. This makes them infinitely flexible and capable of being customized to business needs.

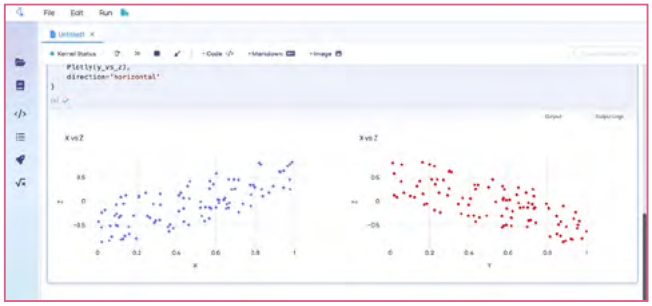

decisionOS comes out of the box with a range of pre-built MetaCells to drive down the time to value.

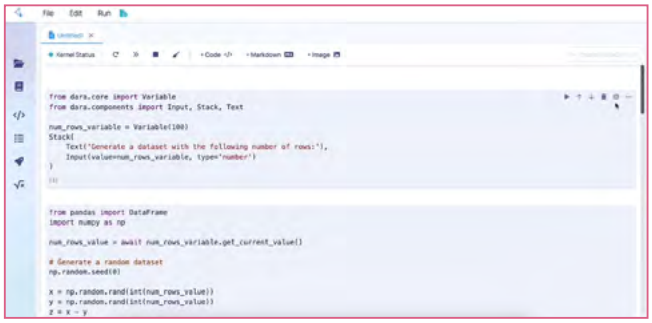

Demonstration of how easy it is to create a MetaCell.

Deployment Options

decisionOS supports a range of flexible deployment options to fulfill your needs:

1. **SaaS Hosted:** decisionOS can be consumed as a SaaS offering with multiple different geographic locations available across the globe. The SaaS offering of decisionOS is HIPAA, ISO27001, and SOC2 compliant, ensuring that your data is protected.
2. **Marketplace:** decisionOS can also be deployed via the Azure marketplace. The advantage of this is that existing cloud credits can be spent in procuring the decisionOS offering. causaLens is actively working with Google and Amazon to enable their marketplace listings, too.
3. **Managed Private SaaS:** By allowing causaLens engineers and deployment processes to manage one of your cloud projects, we can deploy decisionOS as a full SaaS that's only accessible by your users. This option allows you to have all of the capabilities of a SaaS solution, but have it within your own cloud environment.
4. **On-Premise:** decisionOS can also be installed on-premises. This comes at an extra cost, both upfront for installation and ongoing maintenance.

Governance

decisionOS makes data protection easy, bringing enterprise-level security with fine-grained access rights and monitoring for admins or project managers.

Furthermore, decisionOS has numerous compliance certifications, providing users with peace of mind that data protection is covered.

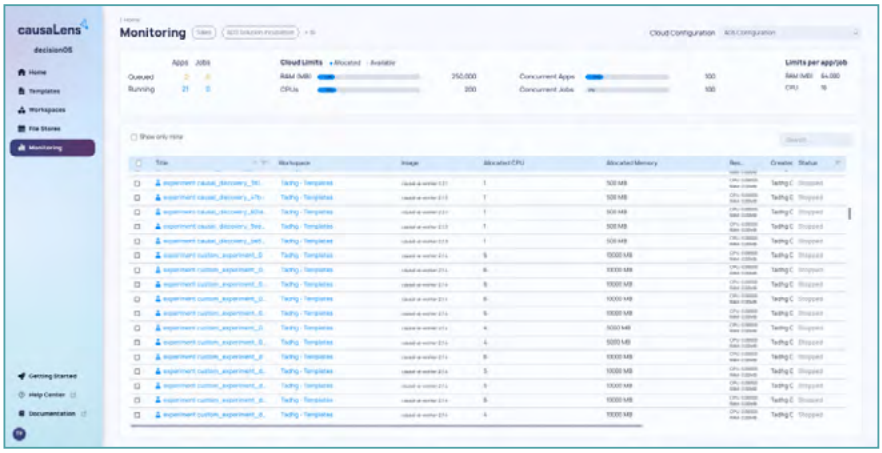

Monitoring panel allowing project managers and administrators to gain an understanding of resource usage rapidly.

Gartner
COOL
VENDOR
2022

Get more info
causaLens.com

info@causaLens.com

causaLens