

SISENSE AI THROUGHOUT

Overview

In today's digital economy, organizations are in a race toward digital transformation, with the fundamental goal of empowering people at every level of the organization to make data-driven decisions and to improve business outcomes. However, despite considerable and growing investment in the field of data and analytics, many organizations still struggle to drive pervasive adoption. In fact, multiple studies show that in most organizations, only a third of individuals are benefiting from existing analytics. Gartner's research has shown that over 65% of organizations have not been able to operationalize their data science projects.

The difficulty in driving adoption and realizing the resulting return on investments is due in part to both the growing complexity of the data as well as scarcity of analytical talent. Additionally, it is critical to understand that analytics can no longer be thought of as an isolated element. It is simply not enough to have a BI application that is distinct and separate from other business applications. While reporting and dashboarding will persist for quite some time, truly transformative organizations are striving to ensure that analytics are pervasive by embedding analytics into every application and ensuring that individuals can turn insights into immediate actions directly from their analytics environments. Furthermore, analytics creates additional revenue engines and business models.

AI and ML: key drivers of digital transformation

Artificial Intelligence (AI) and **Machine Learning (ML)** are key areas of investment for organizations focused on digital transformation. AI helps us remove much of the tedium of trying to figure out how to connect all of this data as well as to help data professionals understand the data to determine how it should best be leveraged. This includes such things as AI-powered data cleansing and modeling as well as data profiling of the underlying data itself.

Machine Learning implementations have been successful in bringing considerable improvements in outcomes across a wide variety of industries, increasing bandwidth for decision-making, as well as responsiveness, accuracy, and consistency of results. Despite soaring demand for and investment in machine learning systems, however, most companies have not yet taken advantage of ML to its fullest potential. There are both resource and experience gaps: not enough data scientists, and a lack of experience in knowing which ML models to deploy for any given problem.

Automated Machine Learning (AutoML) is quickly filling this vacuum by automating as many steps in the machine learning pipeline as possible. Widely hailed as the next wave of machine learning, AutoML both automates the process of applying machine learning to real-world problems while democratizing access to people with limited experience. AutoML thus increases productivity by automating repetitive tasks, keeping valuable data science resources more focused on business problems and less on the models.

Operationalizing AI and ML: our commitment to Augmented Analytics

Artificial Intelligence (AI) capabilities permeate every aspect of the Sisense platform, from the machine learning algorithms that drive a host of functionalities in the core platform, to certified add-ons, to the foundations

underlying our context-based insights. However, we've learned that it isn't enough for an analytics platform to be infused with AI. To support a comprehensive enterprise AI strategy, in addition to our proprietary AI algorithms, we need to support the operationalization of the shelf libraries (AutoML), as well as the ability to apply them on small chunks of data.

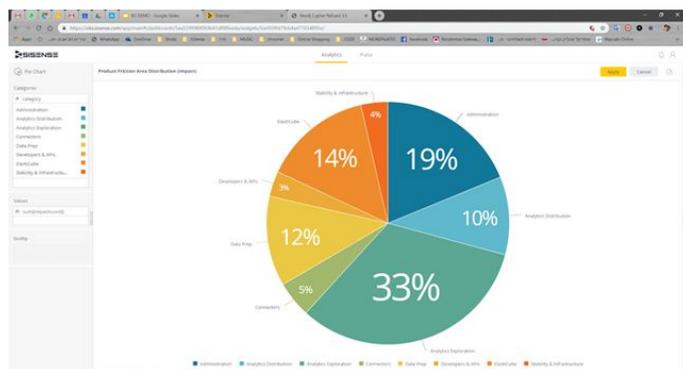
An operational data science pipeline should be able to ingest new data along with continuous support of model improvements that keep the production system stable. Operationalizing AutoML also means handling volumes of training data, model ensemble performance, and dependencies on the external libraries and toolkits. Additionally, AutoML automates the selection of the tools suitable to solve the given problems, as well as handling pre-processing and post-processing steps. While it's difficult to productize ML, Sisense is committed to offering a comprehensive data and analytics platform and tools that empower builders to simplify complex data and deliver insights to everyone inside and outside their organizations.

Augmented Analytics can thus be defined as relieving a company's dependence on data scientists by automating insight generation through ML and AI algorithms, which is at the core of the Sisense mission. When we talk about "AI and ML throughout," we mean that the Sisense data and analytics platform enables everyone across the organization can perform their best work with minimal hassle. Sisense AI and ML capabilities empower end-users with a deeper exploratory environment and more ways to engage with data, leading to faster time to insight and action, coupled with the ability to deploy richer analytics to fewer data-savvy individuals.

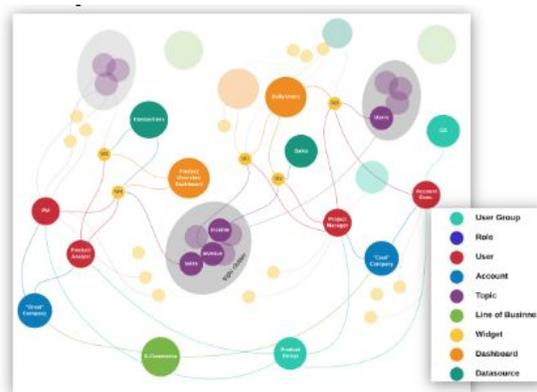
Laying the Foundation with the Sisense Knowledge Graph

Underlying many of the AI and ML investments throughout our platform is our proprietary Knowledge Graph. It 'studies' the associations between all entities on the system (users, data, models, widgets, formulas, clicks, filters, and more) to understand the collective wisdom of the organization leading to more "human" and context-aware outcomes from our AI engine.

Our patent-pending algorithms that power the Knowledge Graph leverage 10 years of anonymized historical data from all Sisense clients' past interactions regarding the usage of users, data, widgets (objects) and dashboards. This jumpstarts the delivery of relevant, powerful and context-based insights from the start. Recommendations that deliver different perspectives and views of a user's KPIs are enhanced by the usage pattern of all users across the system, from data modeling to analytics.



A simple pie chart in Sisense



Behind the scenes, our Knowledge Graph maps out the relationships between the data, users, queries and more.

As discussed, Sisense offers AI- and ML-powered capabilities throughout the platform, serving all the various stakeholders from advanced data teams to the end business user. In the next few pages, we shine a spotlight on Sisense Forecast and Sisense NLQ, two of the most exciting AI and ML developments in the Sisense platform. This is followed by summaries of several additional features and add-ons that make up Sisense AI throughout, which can roughly be broken into three categories: AI-assisted data prep, automated insights, and advanced analytics.

Spotlight on Sisense Forecast and AI Trends: Predicting future outcomes and visualizing trends from historical data

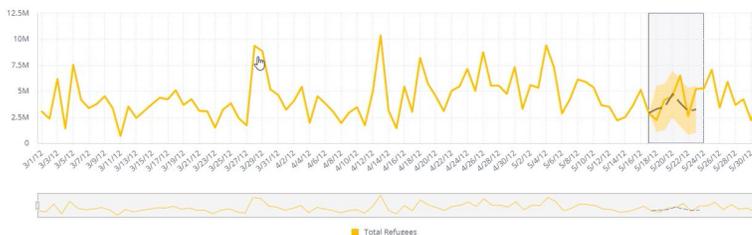
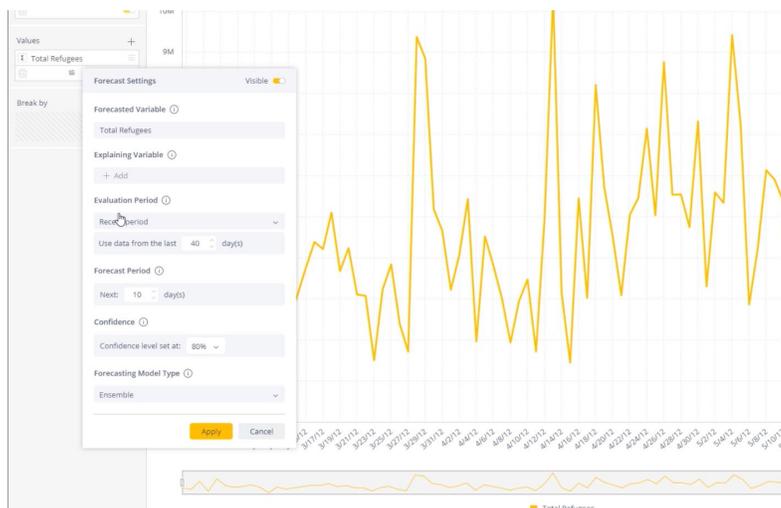
Sisense Forecast and Sisense AI Trends are advanced, ML-powered forecasting and trend functionalities that can be applied via a simple menu option to any visualization based on time-series data including line, area, and column. The two are complementary features; Sisense Forecast allows us to see predictive movements in the future, whereas AI Trends shows what's happened in the past.

With **Sisense Forecast**, Business Analysts can confidently forecast future values and achieve advanced forecasting capabilities using menu-driven, point and click capabilities that allow business users to discern trends and patterns in data immediately.

1). Business analysts can choose from an ensemble of univariate and multivariate Forecasting on any time-series visualization (Line chart, Area chart, Column chart). In the Univariate forecast, the analyst can select a specific forecasting model for the analysis in addition to the ensemble.

2). Sisense's AI engine selects the most effective ensemble configuration, based on the best from the underlying data model.

3). Business users can confidently compare forecast values and actual values for the last period, as well as change the forecast period according to their preferences. Statistical outputs on each forecast model are available from the menu to increase explainability and trust in the forecast.



Sisense Forecast was released in Q4 2019, and is the first analytics service hosted on Sisense Hub. Since then, Sisense has added to the capability, enabling customers to set lower or upper boundaries for their forecasts; adding quick settings for an improved user experience, as well as supporting forecasting models within the Sisense Natural Language Query feature. In the near future, Sisense Forecast will also include the ability to include

or exclude data for specific dates (e.g. if an organization wants to exclude data from the COVID-19 period, this will be possible).

Competing BI tools all require integration with R environments, commercial data science platforms and data science expertise to leverage R forecasting ensembles. Additionally, because we have a pre-processing stage as part of Sisense AutoML, we can run the Forecast on data as small as 30 points.

With **AI Trends**, users can add another layer on top of their visualisations to highlight the trends in the data and make these trends easier to consume. Users can easily draw trend lines using a number of statistical (available locally) and ML-powered algorithms, including exponential smoothing and local estimates (for users who enable cloud access). Compare trends with previous or parallel periods. In addition, AI Trends allows for ignoring anomalies via ML models for users who enable cloud access.

Sisense supports trend lines on all Column, Line, Area, and Bar charts. Sisense also supports the linear trend model within the Sisense Natural Language Query feature. Both Dashboard Designers and Viewers can interact with an “Analyze It” button to apply trend lines, and admins can enable and disable trend types based on cloud models from the Admin page. This functionality was released in Q2 2020 for both Windows and Linux.

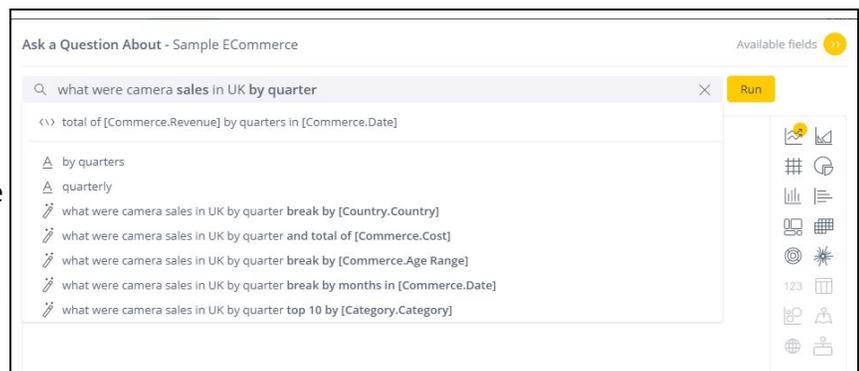


Spotlight on Sisense NLQ: Ask sophisticated questions with natural language

Sisense NLQ (Natural Language Query) enables users to ask sophisticated questions in an easy way with natural language. Users can ask any question, going beyond any predefined widget or visualization created by the dashboard designer. Clicking the 'Simply Ask' button, the user types a question, gets automatic suggestions and immediately sees an appropriate visualization, with full support for spelling mistakes, synonyms and ambiguity.

- Auto-complete:** As the user types a question, the system provides personal suggestions that auto-completes fields and values, filters and break by and analysis over time. These suggestions are constructed using Natural Language Processing (NLP) tools and delivered in context, using Sisense's proprietary Knowledge Graph, which continuously learns from the relationships created between the entities as well as the user's personal preferences.
- Synonyms:** A dictionary of common synonyms is bundled with Sisense NLQ and automatically applied to the NLQ Data Model entities, such as column names, values, formulas, etc. For example, the user can ask "what is my income," and if 'income' is not part of the model but 'revenue' is, the engine will automatically translate 'income' into 'revenue'. When synonyms are applied, the user will see an indication that this occurred and will be able to understand how the system "translated" the term, in our example from 'income' into 'revenue'.
- Spelling:** Spelling mistakes are also presented to the user with suggestions on what to fix.
- Ambiguity:** Sisense NLQ supports ambiguity resolution by providing a ranked list of measure or value appearances in different tables or columns in order to identify the correct measure/value.
- Forecasting and Trends:** Sisense NLQ also supports certain forecast and trends models, meaning that users can ask predictive questions of the data in natural language.

Sisense NLQ includes several ways to aid user engagement and increase adoption. Users can ask questions that go beyond the predefined widgets and visualizations created by dashboard designers. For example, the system displays recent searches and saved searches to enable the user to apply them with a single click. Once the query results are returned, the user can easily modify the visualization by clicking on the desired visualization type.



Furthermore, the user can view how the engine understood his/her question according to standard grammar rules. This way, viewers can modify their text queries if they see that the question was not translated as expected. Viewers may also browse the available fields to understand the categories and topics about which they can ask. Users have the ability to save their search preferences as bookmarks, so that they may be able to easily re-ask their questions (without adding widgets to the dashboard).

Dashboard designers have the ability to enable or disable NLQ and modify the models available.

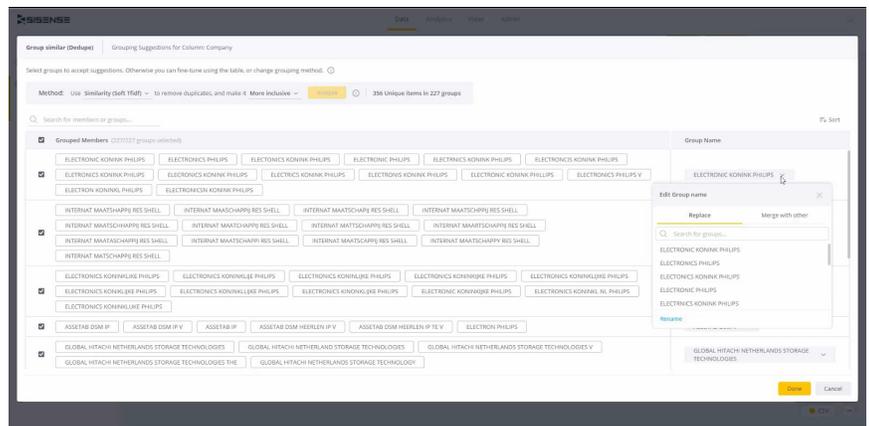
Sisense NLQ was created using the most recent advances in Natural Language Processing, together with the personalization based on Sisense’s proprietary Knowledge Graph. Sisense NLQ was released in Q1 2020, and as of Q3, Sisense Embedded NLQ is also available, making NLQ components accessible in embedded applications.

AI-Assisted Data Prep

The Sisense Elastic Data Engine provides an easy to use environment for working with large disparate data sets. Users can use ready-made point-and-click options or Excel-like functions to transform their data. However, preparing complex and large data sets for analysis can be cumbersome. AI-Assisted Data Prep capabilities within Sisense are focused on reducing the most time-consuming tasks to prepare and model data for analysis. The technical acumen required for data preparation is also lowered, allowing for the expansion of capabilities and responsibilities to more widely skill sets — such as the business analyst.

Group Similar (Dedupe):

The AI-powered Group Similar capability is an augmented data preparation capability that automatically unifies similar values within a given column that are likely the same (i.e. USA, US, United States) and groups them with a few clicks without the need to compose individual SQL expressions. Group Similar has default parameters that generate the model on-the-fly based on similarity functions and sensitivity parameters. The intelligent and intuitive user interface significantly reduces the complexity and time required to group similar values.



Automated Insights

Driving adoption of analytics within organizations is often hindered by a lack of data acumen within the general user population. Users either lack the technical skill or the time to traverse the data and create the content of their own leaving the 'job' of analytics to a relatively small number of overtasked key resources.

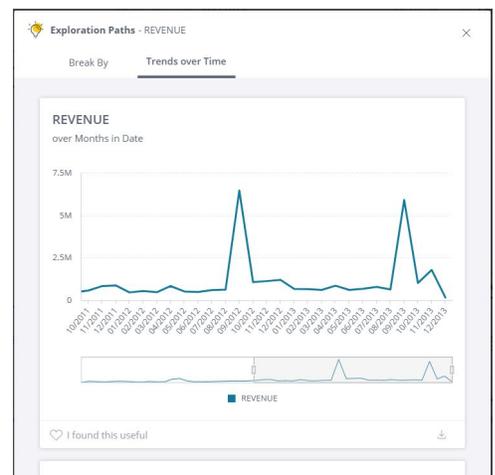
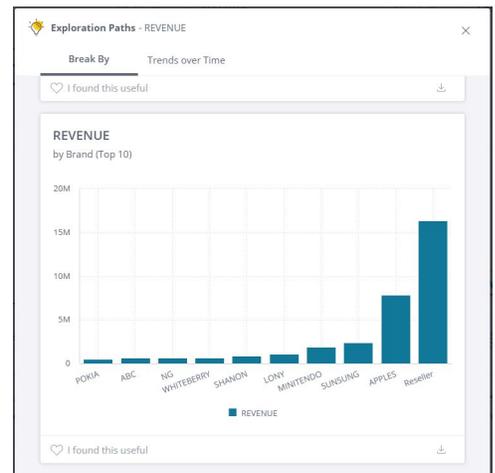
AI can bridge this gap in two ways. First, by automating data analysis and content creation, Sisense AI makes it easier to create meaningful content. And, by providing a guided experience Sisense AI makes it easier for even novice users to gain a deeper understanding of their business.

AI Exploration Paths

Sisense's AI Exploration Paths automatically generates visualizations that explore relationships relating to user KPIs. The AI engine uses a Pointwise Mutual Information based algorithm to automatically identify associations in the data model. Because Sisense leverages the same data models for many use cases, our AI engine is able to learn from all usage patterns from all users of the data, not just those from that particular app.

With Sisense AI-Assisted Exploration, business analysts no longer need to create multiple slices and views of the drivers behind each KPI, making dashboards easier to consume. AI Exploration Paths examines all possible permutations and calculations related to the KPI. Those with statistically relevant correlations automatically generate new visualizations that include break-by and time dimensions so users can analyze the KPI over time.

The resulting visualization can be any number of chart types containing a break-by analysis of the KPI of interest, sliced by a number of correlated dimension fields. The underlying algorithm not only improves from the feedback given by the business users, but also from all end-user activity relating to the data across the entire platform. Results are displayed in an interactive carousel and accept feedback from the end-users regarding the relevance of the insight so that the AI engine continues to learn over time.



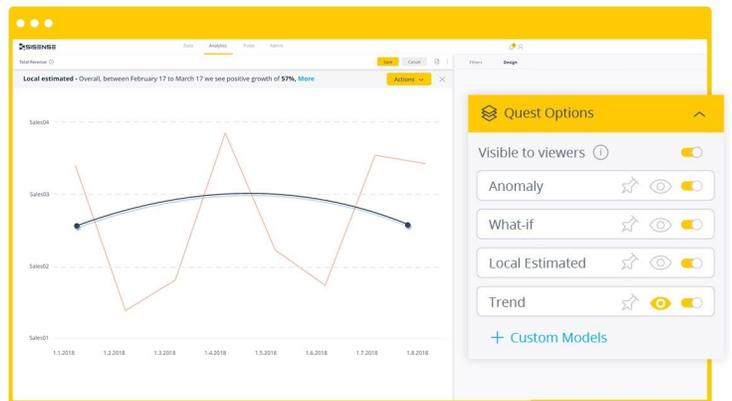
Sisense Quest

Sisense Quest, a certified add-on, harnesses the power of data science for all, enabling builders to deliver advanced analytics through a simple menu to optimize the insight-to-action experience.

Sisense Quest comes with out-of-the-box data science models that can be applied to any time-series data within the Sisense platform. Business analysts control the experience for their business users, suggesting which types of models should be run on different data sets. The pre-built models include what-if, trend, local estimate, anomaly, Pareto, clustering, summary statistics, percentile curves, two-tailed tests and growth rate models. In addition, advanced users can upload their own Python models, thus integrating custom analytics directly into any time-series visualization.

In addition to indicating which models will be exposed to the end-users, analysts can further customize the user experience by associating recommended user actions to each model. These actions can be as simple as downloading the data or sending an email to certain business users, or something more complex, such as sending a command to an external 3rd party system (e.g. stop production or pause a marketing campaign).

Users gain a deeper understanding of their data and, when necessary, can affect change quickly from directly within the app, which drives greater adoption of BI and analytics throughout the organization.



Sisense Pulse Anomaly Detection

Sisense Pulse continuously analyzes KPIs that have been identified by a user using multiple advanced models developed by our internal data science team. When outliers are detected, the system alerts the user automatically via a combination of email, the Sisense Mobile App, or alternative delivery mechanisms like Alexa or smart bulbs. This is accomplished without the need to configure any threshold or rule.

This screenshot shows the 'Add to Pulse' configuration screen. It includes a table with columns 'Name' and 'Value', where 'Total Costs' is listed under both. Below the table, there are sections for 'Filters' (4 active filters), 'Alert Condition' (set to Automatic), and a description: 'The system will monitor this value and notify you when anomalies are detected.' At the bottom, there are 'Basic', 'Add', and 'Cancel' buttons.

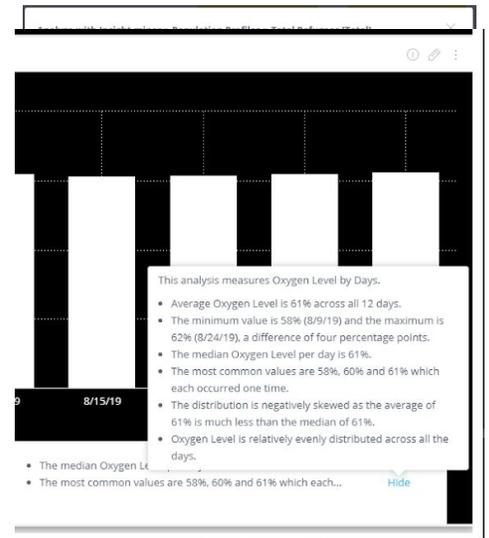
This screenshot shows the notification configuration part of the 'Add to Pulse' screen. It features a 'Message' field for the alert text. Under 'Notifications', there are checkboxes for 'Email', 'Mobile', 'Slack', 'Zapier', and 'Webhook', with 'Email' and 'Mobile' selected. There is also an 'Add Webhook' link. At the bottom, there are 'Basic', 'Add', and 'Cancel' buttons.

Sisense Narratives (Natural Language Generation)

Sisense Narratives leverages natural language generation (NLG) technology to automatically create text-based insights that describe the most important points in a visualization.

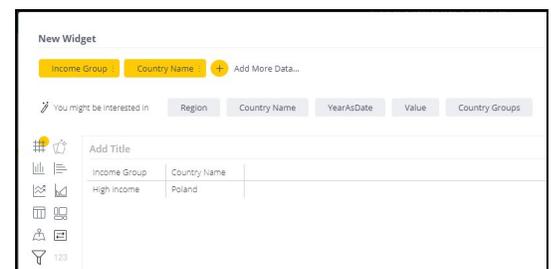
Sisense enables dashboard designers to tailor various aspects of the generated narrative through a menu interface, including the verbosity of the NLG engine on a per-widget level, the sentiment (“up is good” vs “down is good”), the terminology used within the text and the structure (paragraph vs bullet).

NLG is particularly important because it not only helps a user to understand the key points, but it can also help users improve their data literacy as they match the text with what they are seeing in the chart.



Smart Field Recommendations

When designing a new dashboard, Sisense automatically recommends relevant fields from the data model. As the dashboard designer selection fields, our smart field recommendations engine leverages machine learning algorithms to recommend additional fields that may be relevant. As the designer continues to make selections, the system will also recommend the visualization type that is the most appropriate for the data.



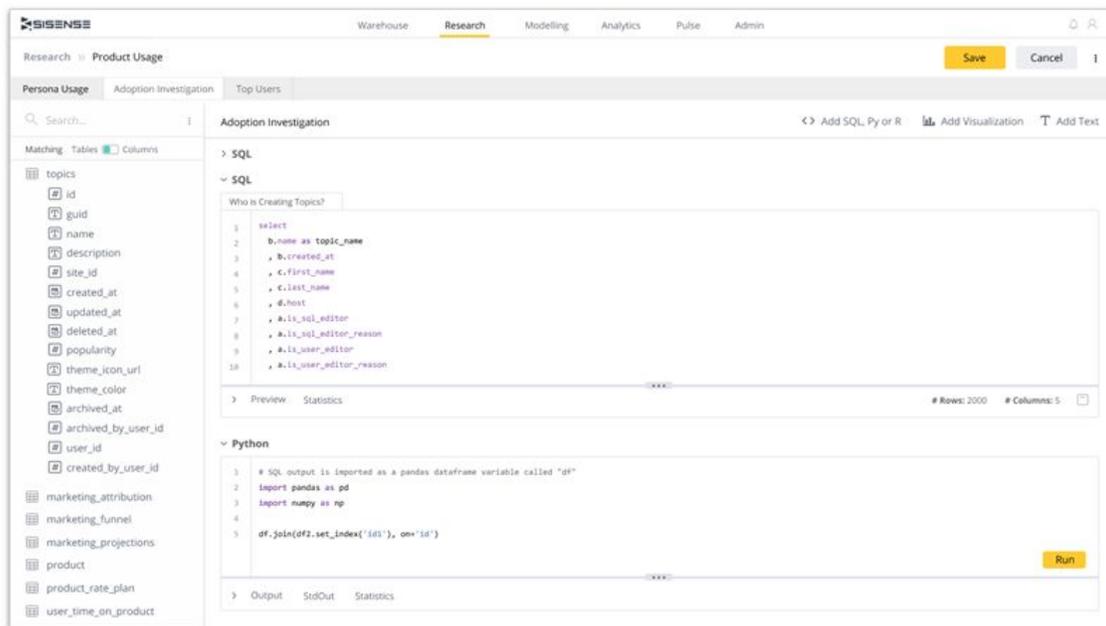
Advanced Analytics

Delivering deeper insights from complex data as part of a digital strategy to drive user adoption is complemented by providing data scientists and advanced data teams with a single platform that combines advanced data preparation capabilities with best-of-breed, open-source machine learning libraries.

Sisense In-Warehouse Data Prep

The Sisense In-Warehouse Data Prep (previously Periscope Data) enables modern data teams to complete more value-based analysis in less time by transforming data in SQL, performing complex statistical analyses using Python or R and operationalize insight across all users in analytical apps - all within the Sisense Platform.

Advanced analytical workflows within one development environment take advantage of a powerful ecosystem of open Python and R libraries - from machine learning to natural language processing to predictive analytics to answer new questions that have yet to be explored. In addition, data from Sisense's Elastic Data Engine can be sent to an AutoML platform such as AWS Comprehend or Sagemaker to leverage bespoke AI services.



The Future is Augmented

For analytics and BI to be truly transformative -- in other words, resulting in significant improvements in business outcomes -- tools need to be smart enough to understand how to connect data sources seamlessly while surfacing meaningful opportunities and insights that are immediately actionable. Augmented Analytics is another way of saying that organizations will need to rely on automated AI and machine learning to make sense of the complex, growing sources and quantities of data while supplementing and enhancing human decision-making.

By investing heavily in AI and Augmented Analytics, Sisense is leading the way in developing technologies and capabilities to fulfill the promise of true digital transformation. Such fulfillment will only be achieved when data-driven decisions are widely adopted across the organization by internal and external stakeholders, and such adoption will only be possible when users will have the tools to understand what their data is showing them and how it can drive greater ROI.

Sisense envisions a powerful new paradigm for its customers in which the user experience is streamlined. The capabilities and underlying technological foundations described in this document point to a workflow that enables users to prepare their data, find patterns in it, and then share and operationalize the findings across their organizations. Backed by Sisense's patent-pending algorithms and proprietary Knowledge Graph, the AI powering these capabilities leverages over a decade of historical data, ensuring relevant, context-based insights from the user's first interaction.

In the new paradigm, time-to-insight is vastly reduced by AI-assisted data preparation; difficult-to-surface patterns are automatically detected and "served" to users of all technical abilities; and advanced data teams can utilize predictive analytics to answer questions that have not yet been asked. With immediately actionable insights embedded directly within business applications, Sisense is driving the era of digital transformation.