ATSCALE

Scaling Data Science and Enterprise AI with a Semantic Layer

Create a Bridge between BI and AI

Table of Contents

The Gap Between AI and BI

This eBook is the fourth in a series focused on the application of semantic layer technology to modern cloud analytics.

The first wave of “big data” technologies focused on managing the petabytes and exabytes of data generated by humans and machines spanning enterprise applications, online user behavior, smart devices, point of sale, lab data, etc... Cloud data platforms made it more practical to consolidate this data within centralized repositories that could be more easily analyzed. Business intelligence (BI) technologies evolved to support analysts and decision makers in making sense of this data. BI focuses on supplying *descriptive* and *diagnostic* analytics to support business users in creating insights that support better decision making.

Data Science and Enterprise AI/ML has emerged over the past several years to include a set of technologies and techniques that leverage artificial intelligence and machine learning models to generate *predictive* and *prescriptive* insights from data. Predictive insights, like sales forecasts, help businesses more effectively plan for the future. Prescriptive insights are suggested actions to optimize a business process.

Despite massive investment in data science and enterprise AI, organizations have struggled to create return due the the complexity of scaling AI models to production and driving systematic business adoption of AI-generated insights.

There is a clear opportunity to address these challenges by leveraging well established BI infrastructure - especially when augmented with a modern semantic layer strategy.

“The collision of once disparate data and analytics markets, fueled by the proliferation of cloud and augmented capabilities, provides the opportunity for data and analytics to drive business driven decisions and outcomes” - Gartner

Challenges to Scaling Enterprise AI

Despite the hype and investment around data science and AI/ML, we are still in the early stage of realizing the impact of enterprise AI. Most companies are just scratching the surface with harnessing the potential of AI and generating real business impact from their investments. Part of the reason is the natural silos that exist between data science teams and the data teams built around BI processes.

For one, data scientists and BI professionals use different tools and technologies. Data science teams spend their day working with Python, open source libraries, containers, APIs, and AutoML platforms. In contrast, business teams typically work with BI platforms like Tableau and Power BI to explore the same data sets.

Data scientists are typically PhDs or computer engineers on separate teams from their BI counterparts. BI teams may sit in IT or within a line-of-business department. Since they operate independently, data scientists are not always aware of or able to leverage investments in data infrastructure built to support mature BI programs.

Data scientists use different terminology than their BI counterparts. BI users talk in the language of measures and dimensions, and focus on descriptive or diagnostic analytics. Data scientists talk about features and predictive model results. Both teams have the same goal – using raw data to model their business in order to generate insights that can help decision-makers be more effective. But the path they take is quite different.

Scaling enterprise AI programs is a key challenge for organizations. Compounding the ever-present shortage of skills, data scientists often waste valuable time wrangling data. Once data teams do generate valuable insights, there is often no consistent mechanism for publishing insights (e.g. predictions, suggestions) back to BI teams and the rest of the organization.

The Silos Separating AI from BI teams

DATA

Data Warehouses, Data Lakes, Apps, Feature Stores, 3rd party Data, etc...

TEAMS

Data Warehouses, Data Lakes, Apps, Feature Stores, 3rd party Data, etc...

TECHNOLOGIES

BI Platforms v. AutoML Platforms, SQL v. Python, etc...

PRIORITIES

Where to invest limited resources? Budget, resources, cycles, attention, etc...

CULTURES

IT Pros v. Business Operators v. Data Scientists, etc...

SEMANTICS

Descriptive v. Predictive v. Prescriptive, Dimensions v. Features, etc...

Bridging the Gap With A Semantic Layer

As discussed in the first eBook of this series, a semantic layer abstracts away the complexity of raw data and creates a platform for teams to create business-focused views of key metrics and analysis dimensions that can be accessed by a broad range of data consumers.

A semantic layer simplifies access to raw data, exposing a business oriented representation of different enterprise data sources including data shared by partners and data acquired from third-party sources.

A semantic layer is where business-oriented logical data models are created and maintained. For BI, this is where metrics and dimensions are defined - tying the views exposed to business users to raw data sources. The modeling environment also enables composability and reusability of models and definitions. The same utilities can be used by data teams to build specific cuts of data that define ML features. Defining complex calculations or setting up time relative features of various lags or windows can be done far easier in a modeling tool than writing complex SQL.

In this way, AtScale can be used to create a set of managed features that can be used for exploratory data analysis, model training, or production models. As these managed features are queried in the form of BI queries, python scripts, or API calls, a semantic layer platform like AtScale serves query traffic to the underlying data platform in the form of optimized SQL. As will be discussed in a later section, this approach to feature serving simplifies data pipeline creation and maintenance.

As predictions are generated by AutoML or AI platforms, they can be written back to data platforms through the semantic layer to inherit the semantic context defined within the model. This approach creates a set of semantic predictions that can be explored from existing BI platforms alongside actuals using the same analysis dimensions. This simplifies business consumption of AI-generated insights and makes it possible to reach much wider audiences.

A semantic layer like AtScale has the potential to reduce the complexity and cost of bringing new models to production as well as generate greater return by encouraging business adoption of AI.

Building Managed Features within the Semantic Layer

Upwards of 80% of a Data Scientist's time is spent preparing data for exploration and model training. Feature engineering is a substantial portion of this data wrangling time. The semantic model can serve as a starting point for building a set of business-vetted features.

The semantic layer is a repository for pre-defined business metrics that have already been defined by the business. This source of "business actuals" for any number of metrics including sales revenue, costs, sales margins, unit sales, shipments, inventory levels can be directly leveraged by data science models with no wrangling or risk of miscalculation.

Furthermore, the semantic layer can be used to quickly build new calculated metrics that can be centrally managed and reused for different ML models. This is particularly useful for time-relative measures including different windows and lags for autoregressive and moving average based time series. Managing time relative metrics within the semantic layer radically simplifies data pipelines, makes sharing and reuse possible, and reduces chance of miscalculations.

It can also be highly advantageous to use governed dimensional hierarchies when engineering features versus reproducing for every model. Standardized dimensions for time, product geographies, and other master data can be shared across different use cases. This can save time and ensures that predictions generated by ML models will be usable and understandable by business users.

Curating Managed Features within the Semantic Layer

Data scientists will draw on a wide range of data sources to build features for exploration and training of new models. For data managed within a data warehouse or lakehouse, the AtScale semantic layer can be used to centrally define managed features and ensure consistency across different use cases.

Furthermore, new features created by AutoML or AI platforms can also become managed features. For instance, an AI platform may suggest an expanded set of time relative features of different windows or lags that improve model performance. AI-Link can write these features back to the model as managed features. These managed features inherit full semantic context, making them more discoverable and easier to work with, consistently, at any stage in ML model development.

The AtScale semantic layer can be used in conjunction with a feature store like FEAST to aggregate the superset of features regardless of where they are sourced. Features can be served directly from AtScale, or through a feature store to train models in AutoML or other AI platforms.

Curating a set of managed features within AtScale, whether or not they are integrated with a feature store, has a few advantages. First, they can be centrally governed for definitional consistency that ensures that business context is retained regardless of where they get used. Second, AtScale provides detailed lineage information for even the most complicated calculated metrics. Finally, managed features can be protected from underlying changes to data - in many cases data consumers will be completely unaware if data stewards redefine the source data used to build a metric.

Integrating AtScale Managed Features with a Feature Store

Sample Python Code

```
project.create_feast_repo(project_name="AtScale_AI",
features=prediction_features,
entities=['item', 'state'],
timestamp='date',
view_name='dataset_ts',
ttl=datetime.timedelta(weeks=0),
force_rewrite=True)
```

Harden ML Pipelines by Serving Features from AtScale

It's important that production AI/ML models have consistent and reliable access to raw data. Changes in the raw data can disrupt model pipelines. This makes it important to insulate production models from changes in the underlying data. There's a need to harden data pipelines for production models so they can deliver feature data accurately, at scale, while meeting performance and recency requirements.

Managed features can serve as this necessary buffer from changes to raw data sources. AtScale exposes the semantic layer to AI workflows with bi-directional Python connectivity, so models can draw directly from managed features that are connected to live data in an underlying cloud data warehouse or lakehouse.

As managed features are requested via AI-Link by a python script, AtScale dynamically pushes optimized queries to the underlying data platform while ensuring high performance. No data is persisted within the AtScale layer - all features are delivered on-demand through virtualized query pipelines with all SQL-based transformations taking place on the data platform.

Governance also becomes important at the stage of feature serving. Real time enforcement of security and access control policies can be managed within the semantic layer. Likewise, governance of query performance and data platform resource consumption can be managed by a semantic layer platform like AtScale.

Make AI-Generated Insights More Accessible with Semantic Predictions

One of the most significant challenges for enterprise AI is driving business adoption of AI generated insights. Oftentimes, the predictions generated by production models get stranded in isolated data science tools or .csv file dumps to a data lake. Deriving value from data science means getting AI-generated insights in front of decision makers quickly and within a familiar analytics experience.

AtScale supports ML model results pipelines that manage the writeback of predictions to the underlying data platform while automatically inheriting the full context of the semantic data model. This makes prediction datasets immediately accessible by business users using existing business intelligence infrastructure.

AtScale uses the concept of Semantic Predictions to represent the idea of prediction data sets with full semantic context. Semantic Predictions can be explored leveraging the dimensional hierarchies to drill up and down on large prediction datasets. Furthermore, predictions can be analyzed alongside business actuals using existing analysis dimensions in the same BI dashboards.

The only way to drive business adoption of enterprise AI is to integrate model-generated insights into existing enterprise analytics workflow. This approach builds on infrastructure and data literacy initiatives already in place to support traditional BI.

Consume Semantic Predictions Alongside Actuals

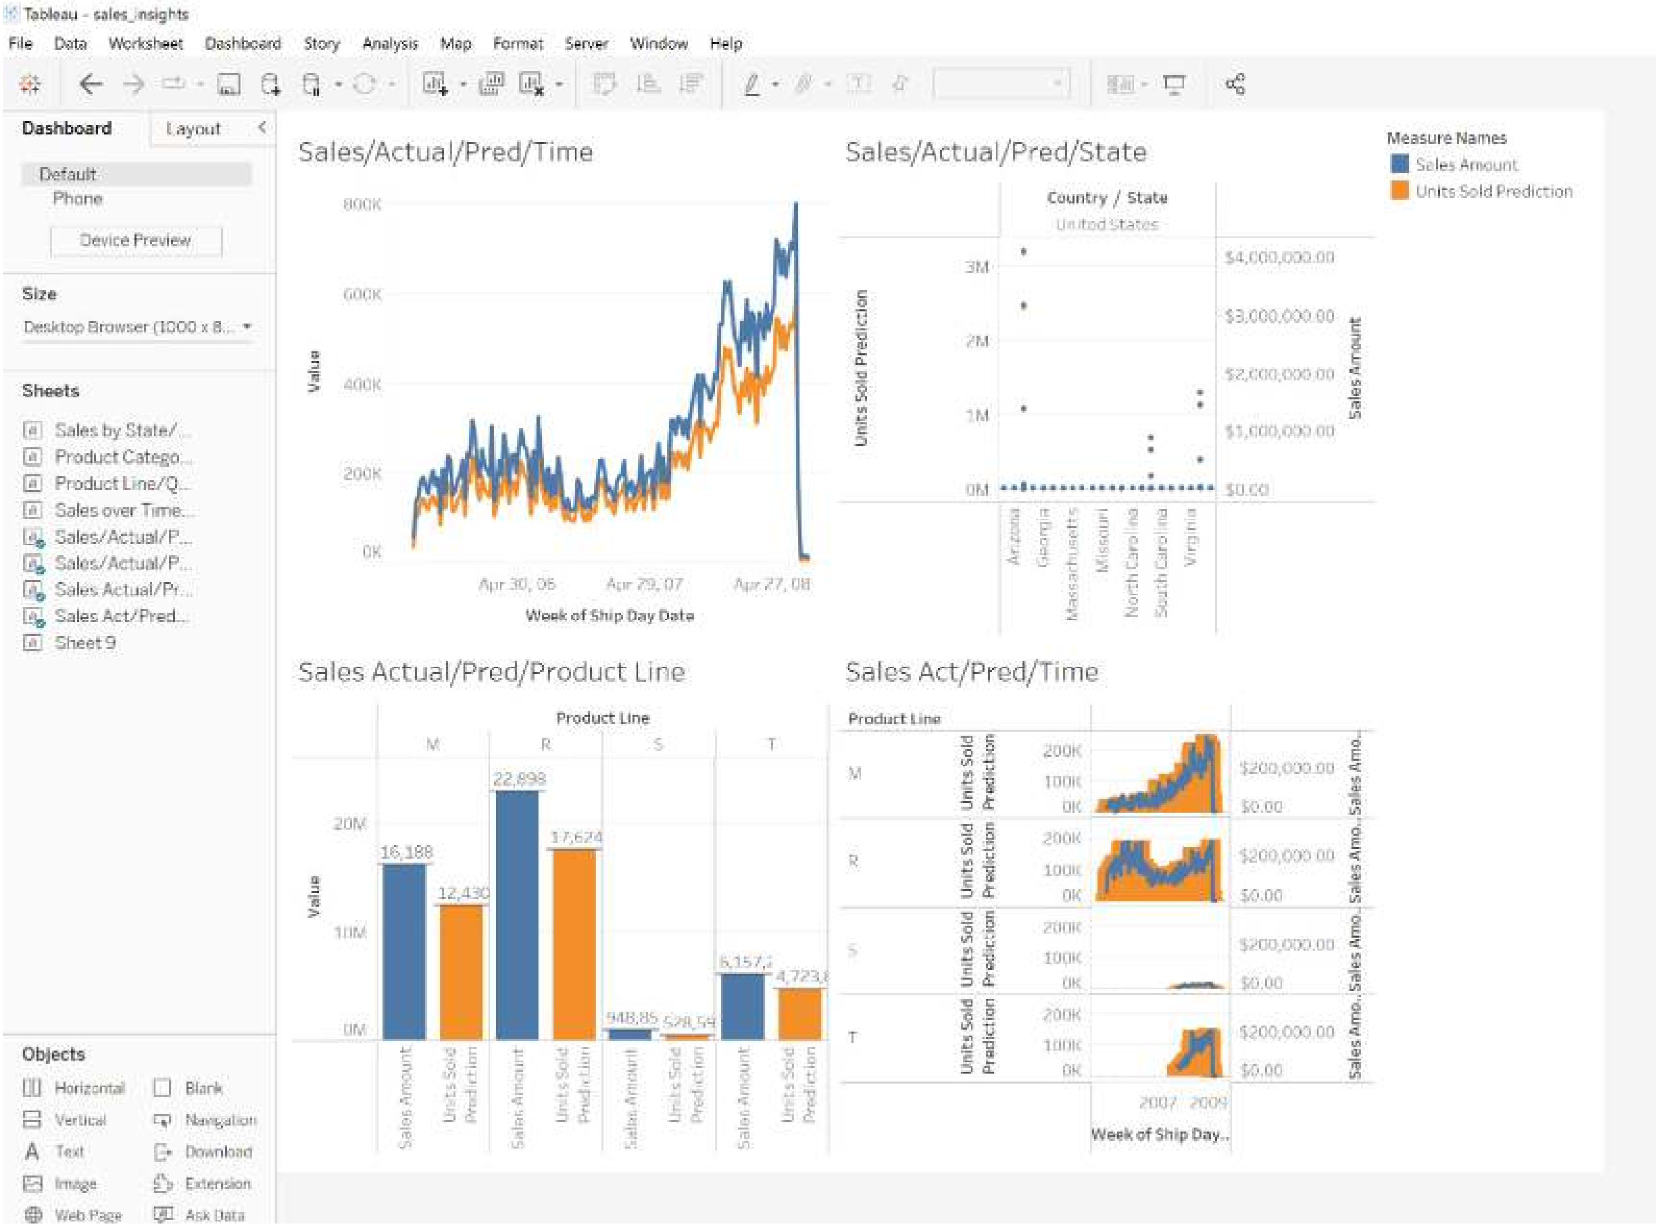

Deep Dive: A focus on Business Forecasting with Predictive Models

We've explored how AtScale democratizes predictive and prescriptive analytics. Let's take a deeper look at how AtScale supports data scientists as they support business forecasting with time series and predictive models, a core activity of enterprise AI.

Time series analysis is based on complex statistical techniques and ML algorithms. Fortunately, the hard work of model selection and model training can be done within a variety of tools such as AutoML platforms. In practice, the real challenge is managing the data used for time series analysis.

Time-relative data is complicated to work with. It involves creating and managing multiple variations of calculations performed on the same metric (e.g. sales). Most commonly, time series data incorporates lag measurements (e.g. comparisons with the same day of the previous week) and window measurements (e.g. three-day running averages). Data sets use different definitions and granularities in time dimension, and don't necessarily come with the aggregation logic (such as hourly or daily) that data scientists need. Models can incorporate data from different systems (such as ERPs and CRMs) and from outside the organization (e.g. an external temperature feed).

Preparing and maintaining clean data across these sources is complicated. Furthermore, maintaining data consistency as models move into production is critical to producing reliable results.

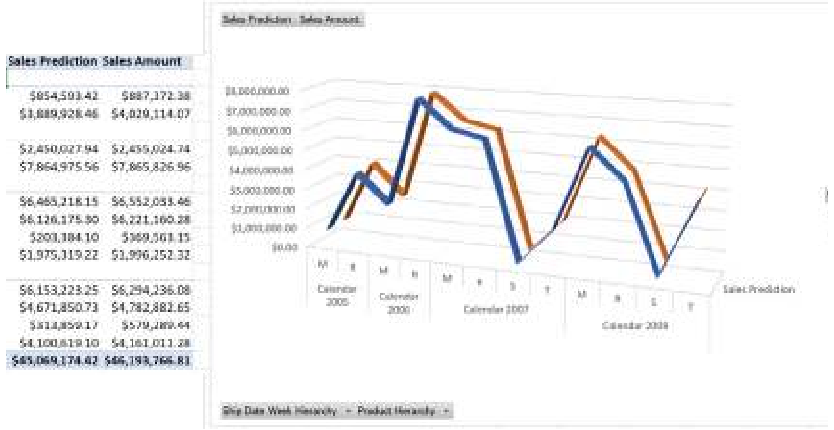

Combining AtScale with a modern cloud data platform like Snowflake or Databricks creates a natural foundation for simplifying data management for time series analysis. AtScale helps data teams create “conformed” time dimensions that map different data sets to a common hierarchical expression of time. This allows data scientists and supporting data teams to quickly and consistently create time-based features across different projects and at different stages of a model development. Time values and time-relative metrics become Managed Features that are “pre-vetted” by the business and available for direct use with a high degree of accuracy.

Cloud data platforms provide the tools for managing large tabular data sets. This data is typically structured for general analytical usage. AtScale lets data scientists and data engineers interact with business-ready forms of raw data with familiar tools like Python notebooks. AtScale manages the translation of the managed feature (e.g. week over week change in sales) to the raw SQL to pull the value from the raw data. This form of feature serving is highly efficient, highly accurate, and resilient to changes to underlying data structure.

AtScale for Bridging AI and BI

To stay competitive, enterprises need to focus on building scalable data science and enterprise AI programs that deliver real business impact. AI enabled analytics have many of the same requirements as traditional BI. Despite the differences discussed in this ebook, the objectives of both AI and BI are the same and modern enterprise analytics programs need to address both.

The real goal of both AI and BI is to deliver descriptive data and AI-augmented insights to business users to support broader insight creation and decision making. This means making data and AI available to users in whatever means and whatever form is most useful. Storing massive amounts of data and generating the most accurate forecasts possible are both completely useless if not adopted and leveraged by business users. Many organizations are struggling to show return on data, data science, and AI technology investments because they are unable to move AI models to production.

A semantic layer platform like AtScale provides a unique set of advantages that helps move more models into production while directly supporting business interaction with data and AI-generated insights. Leverage existing BI infrastructure with a modern semantic layer strategy provides a path toward getting more value from cloud data platforms and from AI investments. AtScale can serve as a control plane for coordinating insight creation across an organization - ensuring data assets are usable by both human users and AI resources as well as ensuring AI-generated insights are accessible by business decision makers where and when they are useful.

Take the Next Step

REQUEST A DEMO →

CONTACT US →

LEARN MORE →

Read The Practical Guide to Using a Semantic Layer

Download The Complete Buyers Guide to a Semantic Layer

Review How does using a semantic layer impact cloud data warehouse performance? (benchmark reports)

Get advice from fellow data & analytics leaders on how to scale smarter data-driven decision-making

Stay current on analytics strategies with short articles and webinars on topics in analytics strategies

About AtScale

AtScale enables smarter decision-making by accelerating the flow of data-driven insights. The company's semantic layer platform simplifies, accelerates, and extends business intelligence and data science capabilities for enterprise customers across all industries. With AtScale, customers are empowered to democratize data, implement self-service BI and build a more agile analytics infrastructure for better, more impactful decision making. For more information, please visit www.atscale.com and follow us on LinkedIn, Twitter or Facebook.