lightup

Anomaly Detection

Lightup provides a rich catalog of Anomaly Detection (AD) algorithms purpose-built to accurately detect data quality issues at scale. These algorithms have been developed using statistical processing and machine learning (ML) techniques — field-tested and refined based on real data from multiple enterprise customers.

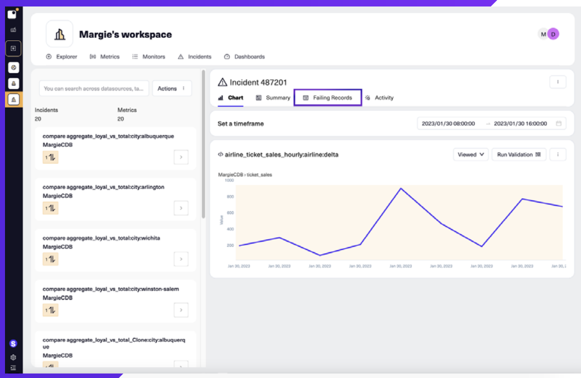

3 Advanced Anomaly Detection Algorithms

Proven very effective in the field, Lightup supports three types of advanced anomaly detection algorithms. Lightup automatically selects the most suitable algorithms for each Data Quality Indicator (DQI) defined in the system, with options for users to customize the Anomaly Detection configurations to meet their desired outcome.

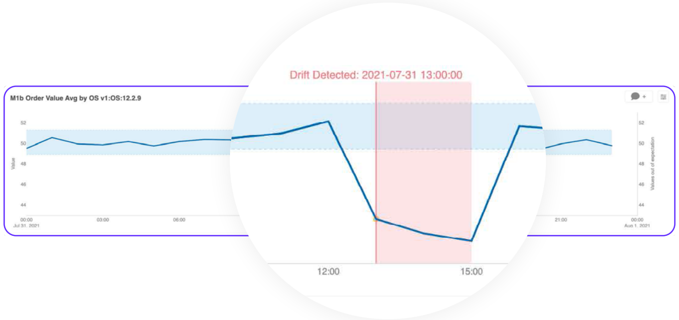

1. Values Outside Expectations

This algorithm detects incidents where a data point does not match expectations predicted from historical patterns. Seasonality and trends observed in the DQI are taken into account when learning expectations from past data, yielding a robust monitor that accurately detects data quality incidents regardless of signal shape complexity.

2. Sharp Change

This algorithm pinpoints incidents where a metric suddenly moves more than expected. The intuition of this algorithm knows data quality issues normally present as sharp deviations from normal DQI behavior. Any seasonality in the signal is taken into account, while ignoring small level changes regarding long-term trends.

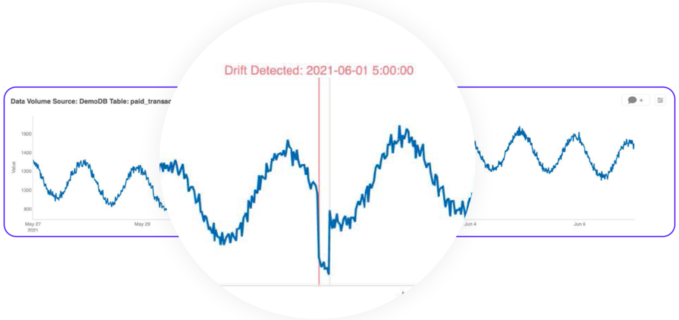

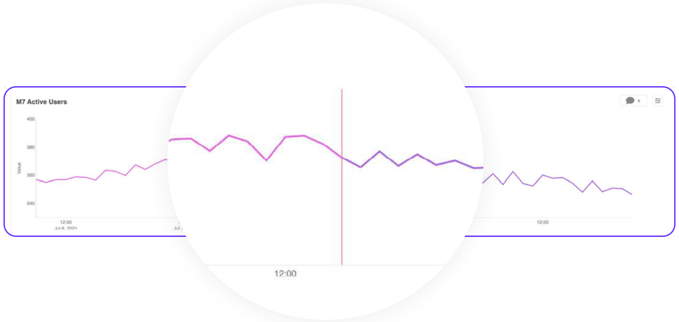

3. Slow Burn Trend Change

This algorithm detects changes in long-term metric trends, very useful for early detection of trend changes that are usually caught too late because of the slow burn nature of such trend changes.

For example, if the number of users on the platform that used to grow at 1% week over week starts decaying 1% week over week, trend change would be the best algorithm to quickly spot this and take action.

Algorithm Training

Enabled by an easy toggle switch click, Lightup's anomaly detection models continuously evaluate the behavior of the relevant metric, flagging outliers or abnormal patterns as incidents.

For optimal performance, anomaly detection models need sufficient training data. Within minutes, these models are trained on the past few weeks of data for the metric, allowing users to provide feedback for fine-tuning, if necessary.

With more complex data elements, users can specify as many unique historical date ranges as needed to make the models seasonally aware and trend accurate.

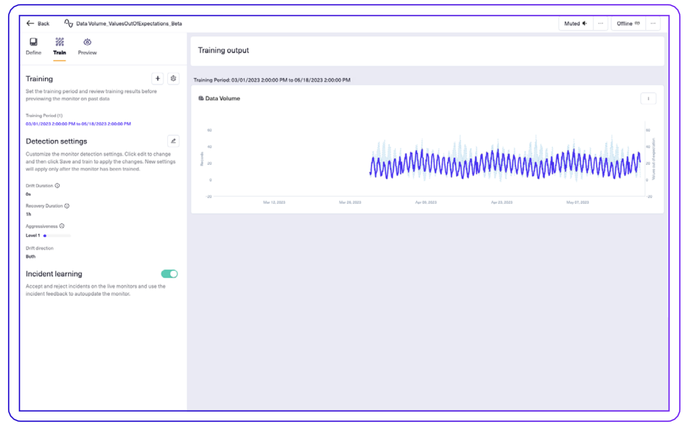

Additional Features for Accurate Detection

In addition to powerful anomaly detection techniques — purpose-built to accurately catch data quality issues — Lightup also provides key features that make the anomaly detection performance easy to supervise, even for new users. Simply put, Lightup provides highly accurate anomaly detection in a no-code/low code environment, accessible to all users with varying levels of statistical and technical expertise.

Rule Preview

Lightup provides an intuitive backtesting/preview workflow that lets users quickly configure detection criteria that works for their requirements, without requiring expertise on the actual statistical models being used.

Each data quality rule can be backtested on historical data using the preview workflow, allowing users to assess rule performance before adding the rule to the live production environment.*

After previewing the results, users can fine-tune the rule using simple settings, such as rule aggressiveness — more aggressive to catch more incidents, less aggressive to catch fewer incidents.

Backtesting allows users to ensure the AD rule will accurately detect the kinds of data quality incidents they want to catch.

**Longer training periods may be needed for more complex seasonality scenarios such as yearly seasonal signals. Most DQIs show one of non-seasonal, daily, or weekly seasonal behavior.*

Online Feedback

Users can pass online feedback to the Lightup system by rejecting data quality incidents that the system cites.

This online feedback is used to:

- Add supervised training to the AD model for improved model accuracy
- Ensure that AD rules continue to be accurate, with low maintenance for end users

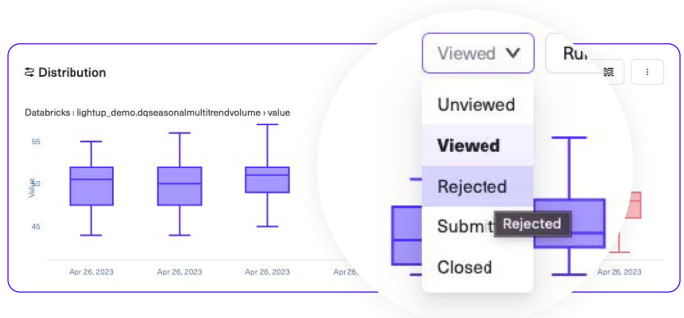

GET IN TOUCH

Info@lightup.ai

lightup.ai