

New Analytics Pages Help

The newer Analytics pages add in-memory analysis for pivots, variance, data quality, ranking, regression, time summaries, outliers, chart recommendations, audit summaries, map readiness, cohorts, funnels, Pareto classes, drift checks, KPI calculations, and data dictionaries. Results can be reviewed, linked back to source records, sent to AI, and exported where supported.

Page	Main input	Main output
Analytics Dashboard	Current report data	Tiles and previews
Pivot	Row, column, value, aggregation	Cross-tab report
Variance	Base, compare, group, value	Change and contribution
Comparison Reports	Periods, groups, queries, or files	Base vs compare
Profiling	Every field	Type/count/blanks/distinct/min/max
Data Quality	Report data	Missing, duplicate, invalid, suspicious rows
Ranking	Category/group and value	Top/bottom/average ranks
Regression	X and Y fields	Equation and predictions
Trends	Equation and X value	Interactive prediction chart
Time Based Summaries	Date and value fields	Day/week/month/quarter/year summaries
Time Series	Date, value, window	Moving averages and rolling totals
Outlier Flagging	Row and value fields, method, threshold	Flagged unusual records
Correlation Threshold	Numeric fields and threshold	Filtered correlation pairs
Chart Recommendations	Category/date/value fields	Recommended charts and dashboards
Audit Summaries	Fields, filters, thresholds, aggregations	Result lineage
Map Readiness	Possible coordinate fields	Map quality and KML readiness
Cohort Analysis	Entity, date, value, period	Cohort retention and activity
Funnel Analysis	Stage, order, value	Drop-off and conversion
ABC Pareto Analysis	Dimension, value, aggregate	A/B/C contribution classes
Data Drift Analysis	Compare field and two segments	Distribution shift
KPI Builder	Dimension, numerator, denominator	Calculated KPI by group
Data Dictionary	Field group and examples	Field documentation and recommended use

Analytics Dashboard

A dashboard is a navigation and preview screen that summarizes available analysis choices for the current report.

Typical screen area

Current report data | Browser size | Selected report | AI | Help | Export

Terms used.

Dashboard = a compact control center with tiles; Tile = one small preview of an analysis page; Current report data = the in-memory rows available for analysis.

Input entries.

The selected report supplies the data. Browser size controls how many tiles fit on each page. AI, Help, and Export buttons open explanation, documentation, or supported output.

Model / algorithm functions.

Builds small preview datasets from current report fields, checks which analytics are meaningful for available field types, and pages tiles responsively.

Meaning of output.

Each tile shows analysis name, miniature preview, and Open link. The preview helps users choose the next analytical page.

Sample result.

Tile	Preview	Open
Pivot	Cross-tab sample	Open
Regression	Equation sample	Open

Shortcut.

Start here when unsure which analysis fits the report.

Input and Fields Selection

- Uses the current report data and available analytics page previews.

Model and Algorithm

- Dashboard tiles summarize analytical pages and provide quick entry points to detailed analysis.

Output

- Tiles open Detail Analytics, Data Readiness, Pivot, Variance, Regression, Quality, Time, Cohort, Funnel, KPI, and other pages.

Why Useful

- Starts analysis from one dashboard and helps users choose the next analytical page.

Suggested Fields

- Use Data Readiness Scanner and the individual analytical pages to choose the best fields.

Pivot / Cross Tab

Builds matrix-style cross-tab reports with row fields, column fields, value fields, and aggregation options.

Typical screen area

Row field | Column field | Value field | Aggregation | AI | Help | Export

Terms used.

Row field = category down the left side; Column field = category across the top; Value field = number summarized; Aggregation = sum, count, average, min, or max.

Input entries.

Choose row and column categories. Choose numeric value for totals/averages, or count when the goal is record counts.

Model / algorithm functions.

Groups records by row/column combination, applies the selected aggregation, and creates dynamic columns for column-field values.

Meaning of output.

The grid shows row labels, dynamic columns, calculated cell values, empty intersections, and totals when available.

Sample result.

Region	Online	Retail	Total
West	12000	8200	20200
East	9400	10100	19500

Shortcut.

Use Pivot when you need matrix totals by two dimensions.

Input and Fields Selection

- Use the current report or imported data after the Search filter is applied.
- Row field selects the category shown down the left side of the cross-tab.
- Column field selects the category expanded into pivot columns across the grid.
- Value field supplies the measure to count, sum, average, min, max, standard deviation, or show as value.
- Aggregation controls how repeated records inside the same row/column intersection are combined.

Model and Algorithm

- Pivot/cross-tab model converts a flat table into a two-dimensional matrix.
- Each source record is assigned to one row bucket and one column bucket.
- The selected value is accumulated inside each bucket using the selected aggregation.
- The page rebuilds the grid so intersections can be compared across row and column categories.

Output

- The first column shows the selected row-field value.
- Generated pivot columns represent distinct values from the selected column field.
- Each cell is the calculated aggregate for that row/column intersection.
- Total shows the row total across generated pivot columns.
- Blank or zero cells mean no matching records or no usable value for that intersection.

Why Useful

- Two category fields can form row and column axes for a cross-tab summary.

Suggested Fields

- Row Field and Column Field dropdowns: select two category/group fields that form the pivot rows and pivot columns
- Value Field dropdown: select the measure to Count, CountDistinct, Sum, Average, Minimum, Maximum, or Standard Deviation

Variance Analysis

Compares base and compare values and explains amount change, percentage change, and contribution to total change.

Typical screen area

Row field | Compare field | Value field | Base/compare selections | AI | Help | Export

Terms used.

Base = starting period/group; Compare = selected comparison period/group; Variance = Compare minus Base; Contribution = row share of total change.

Input entries.

Row field defines groups. Compare field supplies base/compare values such as year, status, region, or scenario. Value field is the numeric amount.

Model / algorithm functions.

Calculates base and compare totals, subtracts variance, divides by base for percent change, and compares row variance to total variance.

Meaning of output.

Base/Compare show the two values. Variance is absolute difference. % Change is relative movement. Contribution identifies drivers.

Sample result.

Group	Base	Compare	Variance	% Change
West	10000	12500	2500	25%
East	8000	7600	-400	-5%

Shortcut.

Click record links to inspect source rows behind base or compare values.

Input and Fields Selection

- Use the current report data after Search filtering.
- Analysis Type selects one of three models: Variance, Percent Change, or Contribution to Total.
- Group Field defines the output rows, such as category, period, location, department, customer, product, or another dimension.
- Value Field selects the numeric measure to summarize, and Aggregation defines how repeated records are combined.
- For Variance and Percent Change, Compare Field supplies the list of Base Value and Compare Value choices; those two selected values must be different.
- For Contribution to Total, Base Value and Compare Value are not used; the page calculates each group share of the selected value total.

Model and Algorithm

- Variance compares a base result against a comparison result by group.
- Percent Change uses the same base/compare groups but expresses the difference relative to Base.
- Contribution to Total ignores base/compare selections and calculates each group value divided by the total value.
- Records are grouped by the selected Group Field and summarized with the selected Aggregation.
- For Variance and Percent Change, Variance is Compare minus Base, and Percent Change is calculated against Base when Base is not zero.

Output

- For Variance and Percent Change analysis, the grid shows the selected group field, Base value, Compare value, Variance, and Percent Change.
- Base and Compare column captions include the selected base and compare values.
- Variance is Compare minus Base; positive values mean compare is higher and negative values mean compare is lower.
- Percent Change shows the relative movement from Base when Base is not zero.
- For Contribution analysis, the grid instead shows the selected group field, Value, and Contribution to Total.

Why Useful

- Compare values across groups, periods, or categories.

Suggested Fields

- Group Field dropdown: select the category, period, location, department, customer, product, or other dimension that becomes each output row
- Value Field and Aggregation dropdowns: select the numeric measure and how records should be summarized before comparison

Comparison Reports

Compares two periods, groups, locations, SQL queries, or imported files using the same group/value logic.

Typical screen area

Comparison type | Base selector/query/file | Compare selector/query/file | Group/value fields | AI | Help | Export

Terms used.

Base set = first dataset/filter; Compare set = second dataset/filter; Match key = group field used to align results.

Input entries.

Comparison type controls where data comes from. Two Queries uses two SQL statements. Two Imported Files uses two comparable files.

Model / algorithm functions.

Builds base and compare tables, groups both by selected group field, aligns matching groups, and calculates counts, totals, variance, and percent change.

Meaning of output.

Base/Compare Records are row counts. Base/Compare Values are aggregated numbers. Links open corresponding filtered records when available.

Sample result.

Group	Base Records	Compare Records	Variance
Product A	45	51	1200
Product B	28	20	-850

Shortcut.

Use comparable columns in both imported files so results can be matched reliably.

Input and Fields Selection

- Choose whether to compare two periods, groups, locations, SQL queries, or imported files.
- Row field identifies the common dimension used to align base and compare results.
- Value field and aggregation define the measure being compared.
- Base and Compare controls select the two sides of the comparison.
- For Two Queries, enter each SQL statement; for Two Imported Files, choose the two files to load and compare.

Model and Algorithm

- Two-source comparison model aligns two datasets by the selected row dimension.
- The page builds separate base and compare datasets, then groups and aggregates each side.
- Matching row values are joined so differences can be calculated side by side.
- Variance is Compare minus Base; Percent Change is calculated from Base where possible.
- Base/compare filters are registered so each result can drill back to its source records.

Output

- Comparison Type identifies the selected comparison method.
- The row-dimension column is named from the selected Row Field and shows the value used to match base and compare rows.
- Base and Compare columns include the selected base and compare values in their captions.
- Variance and Percent Change show the comparison result.
- Base Records and Compare Records link to the exact records used for each side.

Why Useful

- Compare two periods, groups, locations, queries, or imported files.

Suggested Fields

- Comparison Type dropdown: choose Periods, Groups, Locations, Two Queries, or Two Imported Files; then fill the matching base/compare query, file, period, group, or location controls
- Value Field / Aggregation controls: select the numeric measure and summary calculation used to compare base and compare results

Data Profiling

Profiles every field in the report or imported dataset to reveal type, completeness, uniqueness, and numeric range.

Typical screen area

Current report data or imported data | AI | Help | Export

Terms used.

Profile = field-level summary; Blank = missing value; Distinct = count of different values; Standard deviation = numeric spread around average.

Input entries.

Uses the current report dataset or imported data. Search/filter options can narrow the field list.

Model / algorithm functions.

Scans every column, detects likely data type, counts records/blanks/distinct values, and calculates min, max, average, and standard deviation where applicable.

Meaning of output.

Each row describes one field. Blanks show completeness issues. Distinct values reveal IDs/categories. Numeric statistics show range and spread.

Sample result.

Field	Type	Count	Blanks	Distinct
Sales	Number	1000	0	923
OrderDate	Date	1000	2	356

Shortcut.

Use profiling before deeper analytics to choose good category, date, and value fields.

Input and Fields Selection

- Profiles every field in the current report or imported dataset.
- Search limits the records scanned when the user wants profiling for a subset.
- No field selection is required; the page automatically inspects all columns.
- Numeric-looking, date-looking, and text fields are handled differently so each receives meaningful statistics.

Model and Algorithm

- Automatic profiling model describes the structure and behavior of each field.
- The page counts records, blanks, and distinct values for every column.
- It detects likely data type from actual values and column behavior.
- Numeric fields receive min, max, average, and standard deviation where calculation is valid.
- Date fields receive date-oriented min/max values; text fields receive distinct/blank patterns.

Output

- Each grid row represents one field in the dataset.
- Data Type explains how the field is interpreted.
- Count, Blanks, and Distinct Values show completeness and uniqueness.
- Min, Max, Average, and Standard Deviation are populated where applicable.
- Blank statistic cells mean the field does not support that calculation.

Why Useful

- Detect type, blanks, distinct values, min, max, average, and standard deviation.

Suggested Fields

- Automatic profiling: all fields are scanned; numeric fields receive min/max/average/stdev, date fields receive date ranges, and text/category fields receive blanks, distinct counts, and examples

Data Quality

Finds common data-quality problems that can distort reports, analytics, maps, charts, and exports.

Typical screen area

Current report data | Optional selected fields/rules | AI | Help | Export

Terms used.

Duplicate = repeated record/key; Invalid date = not usable as date; Suspicious text = unusually long, mixed, or inconsistent text.

Input entries.

Uses report data and optional field/rule choices. Numeric ranges identify business-limit problems. Category checks find spelling/casing inconsistency.

Model / algorithm functions.

Runs checks for blanks, duplicates, invalid dates, out-of-range numbers, inconsistent categories, and suspicious text patterns.

Meaning of output.

Count shows records failing a check. Severity indicates likely impact. Count links open affected records for review.

Sample result.

Check	Field	Count	Severity
Missing values	Customer	12	Medium
Invalid dates	ShipDate	3	High

Shortcut.

Review high-severity rows before relying on statistics or models.

Input and Fields Selection

- Checks all available fields in the current report data.
- Search filters the records before quality rules are applied.
- Threshold settings control numeric sensitivity for suspicious values.
- Date-like fields are checked for invalid dates; numeric fields for range issues; text/category fields for blanks, inconsistent values, and suspicious content.

Model and Algorithm

- Data quality model applies rule-based checks to identify records that may need correction.
- The page checks missing values, duplicate rows, invalid dates, out-of-range numbers, inconsistent categories, and suspicious text.
- Each issue is stored with a filter pointing back to affected records.
- Duplicate detection compares complete row signatures.
- Category checks look for casing, spacing, and low-quality text patterns.

Output

- Check names the quality check that found the issue.
- Field identifies the affected column when the issue is field-specific.
- Issue describes the problem found in the data.
- Count shows how many records are affected and links to those records when a record filter is available.
- Details gives the extra context needed to review or correct the issue.

Why Useful

- Missing values and duplicate records are counted for the current data.

Suggested Fields

- Date checks: review date-like fields for invalid dates, impossible dates, and missing date values
- Numeric checks: review numeric measure fields for out-of-range values, suspicious extremes, and standard-deviation exceptions
- Category/Text checks: review text fields for blanks, inconsistent spelling/casing, duplicate-looking categories, and suspicious text values

Ranking Analysis

Ranks categories, customers, products, departments, locations, or other dimensions by top, bottom, or average values.

Typical screen area

Category/group field | Value field | Rank type | Group field optional | AI | Help | Export

Terms used.

Rank = ordered position; Top = largest; Bottom = smallest; Average = typical value; Group Value = rank value inside optional group.

Input entries.

Category/group defines what is ranked. Value is the numeric measure. Rank type controls largest, smallest, or average emphasis.

Model / algorithm functions.

Groups records, calculates selected rank measure, sorts according to rank type, optionally ranks inside each group, and builds record links.

Meaning of output.

Rank shows position. Category identifies item. Value columns show Top, Bottom, or Average Value. Records link to source rows.

Sample result.

Rank	Category	Top Value	Records
1	West	124000	88
2	East	118000	79

Shortcut.

Use Average to compare typical behavior instead of volume.

Input and Fields Selection

- Rank field selects the dimension to rank, such as category, customer, product, department, or location.
- Optional Group field creates separate rankings inside each group.
- Value field selects the numeric measure to rank.
- Rank Type chooses Top, Bottom, or Average analysis.
- Top Count controls how many ranked rows are returned; Search filters source records before ranking.

Model and Algorithm

- Ranking model orders dimensions by calculated business value.
- Records are grouped by Rank field and optional Group field.
- The selected Value field is aggregated, then sorted according to Rank Type.
- Top and Bottom use the largest or smallest totals; Average uses average value per record.
- When a group is selected, the page also calculates the rank value within each group.

Output

- Group shows the selected group value when a group field is used; otherwise it is blank.
- Top Value, Bottom Value, or Average Value appears when a group field is selected and summarizes the selected Rank Type for that group.
- Rank shows row order inside the group or complete result.
- Item is the selected rank-field value, and the value column shows the calculated Top, Bottom, or Average result.
- Records links open the rows behind each ranking result.

Why Useful

- Category fields and numeric values can be ranked by top, bottom, or average.

Suggested Fields

- Group Field dropdown: optionally split rankings by category, customer, product, department, location, period, or another dimension
- Value Field dropdown: select the numeric measure to rank, such as sales, revenue, quantity, amount, cost, profit, score, or duration

Regression Analysis

Explains and predicts how one value changes when another value changes, including nonlinear and logistic options.

Typical screen area

X field(s) | Y field | Equation type | Predict Y when X is | AI | Help | Export

Terms used.

X = independent/input value; Y = dependent/output value; Prediction = estimated Y; R-squared = fit strength; Logistic probability = yes/no likelihood from 0 to 1.

Input entries.

X fields provide explanatory values. Y is predicted value. Equation type controls model shape. Predict Y when X is supplies the X used for prediction.

Model / algorithm functions.

Fits selected equation to numeric data, evaluates fit, removes zero coefficient terms, calculates predicted Y, and prepares Trends links.

Meaning of output.

Equation describes relationship. Predicted Y is estimate for chosen X. Fit notes judge reliability. Records links show rows used.

Sample result.

Group	Equation	Predicted Y	Records
All	$Y = 4.851 * \text{pow}(X, 0.8333)$	8.64	120
Yes/No	$P = 1/(1+\text{exp}(\dots))$	0.72	95

Shortcut.

Use Trends to inspect whether the fitted curve makes business sense.

Input and Fields Selection

- X Field is the driver or independent numeric field.
- Y Field is the result or dependent numeric field to explain or predict.
- Optional Group field fits separate models for different categories.
- Equation Type selects linear, polynomial, exponential, power, logarithmic, or logistic-style fitting where data supports it.
- Predict Y when X is supplies an optional X value used to calculate a forecast.

Model and Algorithm

- Regression model estimates how Y changes when X changes.
- The page collects valid numeric X/Y pairs, excluding blank or nonnumeric values.
- For each group, it fits the selected equation type and calculates coefficients.
- The best supported equation is shown without zero-value terms.
- Prediction is calculated by substituting the requested X value into the equation.

Output

- Group identifies the combined segment when grouping is used.
- Equation shows the fitted relationship between X and Y.
- Predicted Y shows the forecast for the requested X value when provided.
- Records links open the rows used to fit each equation.
- Trends links open an interactive chart where the equation can be explored and exported.

Why Useful

- Two or more numeric fields can be tested for prediction and fitted equations.

Suggested Fields

- X Field dropdown: select the numeric driver or independent variable that may explain changes
- Y Field dropdown: select the numeric result or dependent variable to explain or predict

Trends

Shows an interactive equation chart with selected X and predicted Y values, including nonlinear functions.

Typical screen area

Equation | X value | Optional subtitle values from Regression | AI | Help | Export

Terms used.

Equation = formula for Y; Active X = selected input; Curve = generated Y values; Zoom = change visible range without changing equation.

Input entries.

Equation can include pow, exp, ln, and log. X value controls highlighted prediction. Subtitle identifies group/fields from Regression.

Model / algorithm functions.

Parses equation, samples X values across chart range, calculates finite Y values, draws curve, and updates active point on click/scroll.

Meaning of output.

Chart shows X/Y relationship. Highlighted point shows prediction. Exports capture chart. Record links appear when source rows exist.

Sample result.

Equation	X	Y
$Y = 10 + 2 * X * X$	2	18
$Y = 4 + 2 * e^X$	1	9.44

Shortcut.

Use zoom and scroll bars to inspect local detail or wider trend shape.

Input and Fields Selection

- Equation textbox accepts the equation to draw, and X Value supplies the active X used for the highlighted chart point.
- The page can also receive equation, X field, Y field, and group information from Regression or Open Trends.

Model and Algorithm

- The page evaluates the equation over an adjusted X/Y range, supports functions such as log, ln, exp, powers, and arithmetic, and recalculates the active point when the chart is clicked.
- Zoom and scroll controls change the visible range while the chart area stays the same size.

Output

- The chart shows the equation curve, axes, active point, prediction text, and export buttons for Excel and PDF.

Why Useful

- Two or more numeric fields can be tested for prediction and fitted equations.

Suggested Fields

- X Field dropdown: select the numeric driver or independent variable that may explain changes
- Y Field dropdown: select the numeric result or dependent variable to explain or predict

Time Based Summaries

Summarizes values by day, week, month, quarter, or year when date fields exist.

Typical screen area

Date field | Value field | Date aggregation | AI | Help | Export

Terms used.

Date aggregation = grouping dates into periods; Quarter = three-month period; Period record link = source rows inside that time bucket.

Input entries.

Date field must contain usable dates. Value field supplies numeric amount. Date aggregation sets daily, weekly, monthly, quarterly, or yearly grain.

Model / algorithm functions.

Converts each date to selected period, groups records by period, and calculates record count, total, average, minimum, and maximum.

Meaning of output.

Period identifies time bucket. Records is row count. Total/Average summarize value. Links open records for each period.

Sample result.

Period	Records	Total	Average
2026-Q1	220	450000	2045
2026-Q2	180	390000	2167

Shortcut.

Use quarter/year for management views and day/week for operational detail.

Input and Fields Selection

- Date Field selects the date column used to create time periods.
- Date Aggregation chooses Day, Week, Month, Quarter, or Year.
- Value Field selects the measure summarized inside each period.
- Aggregation defines Count, Sum, Average, Min, Max, or other available calculation.
- Search filters the records before period totals are created.

Model and Algorithm

- Period-summary model converts raw date records into grouped time buckets.
- Dates are normalized to the selected aggregation level.
- Records inside each period are grouped together.
- The selected value field is summarized using the selected aggregation.
- A filter is registered for each period so totals can be traced back to source records.

Output

- Period shows the calculated day, week, month, quarter, or year bucket.
- Period Start shows the normalized beginning date for that bucket.
- Records shows how many rows are included in that period and links to them.
- The calculated value column is named from the selected Aggregation and Value Field, such as Sum of Sales or Average of Amount.
- The grid is useful for seasonality, monthly summaries, annual rollups, and period comparisons.

Why Useful

- Date and numeric fields support summaries by day, week, month, quarter, and year.

Suggested Fields

- Date Field dropdown: select the date column used to create day, week, month, quarter, or year periods
- Value Field dropdown: select the numeric measure summarized in each period

Time Series

Calculates moving averages and rolling totals for time-series style reports.

Typical screen area

Date field | Value field | Date aggregation | Number of time periods | AI | Help | Export

Terms used.

Moving average = average over recent periods; Rolling total = sum over recent periods; Window = number of periods included.

Input entries.

Date orders the series. Value is summarized. Date aggregation sets grain. Number of time periods controls moving/rolling window size.

Model / algorithm functions.

Groups values by period, sorts chronologically, and calculates value, moving average, and rolling total for each period using the window.

Meaning of output.

Value is period amount. Moving Avg smooths noise. Rolling Total shows accumulated recent value. Links open period rows.

Sample result.

Period	Value	Moving Avg	Rolling Total
Jan	100	100	100
Feb	120	110	220

Shortcut.

Increase window for smoother trends; decrease it for faster reaction.

Input and Fields Selection

- Date Field selects the chronological field used to order periods.
- Date Aggregation controls whether periods are day, week, month, quarter, or year.
- Value Field selects the numeric measure analyzed over time.
- Number of time periods sets the rolling window size.
- Calculation Type chooses moving average or rolling total; Search filters the source records first.

Model and Algorithm

- Time-series model analyzes ordered period values.
- Records are grouped into periods and sorted chronologically.
- Period totals are calculated from the selected Value Field.
- The rolling window looks back across the selected number of periods.
- Moving average divides the rolling sum by available periods; rolling total keeps the sum.

Output

- Date shows the ordered time bucket.
- Records links open source rows in that period.
- Value shows the summarized value for the period before rolling calculations.
- Moving Average and Rolling Total show the rolling results for the selected number of time periods.
- This helps smooth noisy values and reveal trend direction over time.

Why Useful

- Date and value fields support moving averages and rolling totals.

Suggested Fields

- Date Field dropdown: select the chronological date column used to order periods
- Value Field dropdown: select the numeric measure used for period value, moving average, and rolling total

Outlier Flagging

Flags unusual records or groups based on standard deviation, percentage difference, or business rules.

Typical screen area

Row field | Value field | Outlier method | Threshold | AI | Help | Export

Terms used.

Outlier = unusually high/low value; Threshold = flagging limit; Standard deviation method = far from average; Business rule = user-defined limit.

Input entries.

Row identifies records/groups. Value is numeric measure. Method chooses statistical or rule-based detection. Threshold controls sensitivity.

Model / algorithm functions.

Calculates expected range, compares each value to range or rule, assigns high/low notes, and prepares source-record links.

Meaning of output.

Row identifies affected item. Value is observed amount. Method shows why checked. Outlier Note explains high/low or rule violation.

Sample result.

Row	Value	Method	Outlier Note
Order 1001	9999	Std Dev	High outlier
Order 1044	1	Percent Diff	Low outlier

Shortcut.

Outliers are signals for review, not automatic errors.

Input and Fields Selection

- Row field identifies the record, category, or entity shown in the output.
- Value field selects the numeric measure being tested.
- Rule Type chooses standard deviation, percent difference, or business-rule style checks.
- Threshold controls how sensitive the flagging should be.
- Search filters records before outlier statistics are calculated.

Model and Algorithm

- Outlier model flags values that are unusual compared with the selected rule.
- For standard deviation rules, the page calculates average and standard deviation.
- For percent-difference rules, the page compares values against a baseline.
- Business-rule thresholds flag values above or below configured limits.
- Each flagged row receives a reason explaining the rule that matched.

Output

- Row is the source-record number and is linked to the matching record in Data Explorer.
- Field shows the selected numeric field being checked.
- Value shows the flagged value from that row.
- Method and Reason explain which outlier rule was applied and why the row was flagged.
- Average and Std Dev show the baseline statistics used by standard-deviation and percent-difference methods.

Why Useful

- Numeric values can be checked for unusual deviations or business-rule exceptions.

Suggested Fields

- Row Field dropdown: select the row label, category, entity, or record identifier shown for each flagged outlier
- Value Field dropdown: select the numeric measure tested by standard deviation, percent difference, or business rule limits

Correlation Threshold

Filters correlations to focus only on numeric field pairs above or below a selected threshold.

Typical screen area

Numeric fields | Threshold | Direction/filter | AI | Help | Export

Terms used.

Correlation = strength/direction of numeric relationship; Positive = move together; Negative = move opposite; Threshold = minimum strength to show.

Input entries.

Numeric fields define possible pairs. Threshold controls required strength. Direction/filter focuses on positive, negative, high, or low relationships.

Model / algorithm functions.

Calculates correlations for numeric field pairs, compares values to threshold, labels strength, and hides weak/noisy pairs.

Meaning of output.

Field 1/2 identify pair. Correlation ranges -1 to 1. Strength describes size. Note helps interpret direction.

Sample result.

Field 1	Field 2	Correlation	Strength
Sales	Profit	0.88	Strong
Discount	Margin	-0.63	Moderate negative

Shortcut.

Use threshold filtering when the full correlation grid is too noisy.

Input and Fields Selection

- Uses eligible numeric fields from the current report data.
- Threshold Direction controls whether to show values above, below, or matching the selected strength rule.
- Threshold Value sets the minimum or maximum correlation strength to display.
- Search filters field-pair results after calculation.
- Avoid ID/index fields because they can create misleading correlations.

Model and Algorithm

- Correlation-threshold model filters numeric relationships by strength.
- The page calculates pairwise correlation between numeric field pairs.
- Each correlation is compared with the selected threshold rule.
- The result is classified by direction and strength.
- Only pairs meeting the threshold are shown so the user can focus on stronger relationships.

Output

- Field 1 and Field 2 identify the numeric pair being compared.
- Correlation ranges from -1 to 1.
- Strength shows Moderate or Strong based on the absolute correlation value.
- View shows Positive when the fields move together and Negative when they move in opposite directions.
- Only field pairs that pass the selected threshold and view filter are displayed.

Why Useful

- Filter correlation pairs by minimum strength and focus on the strongest relationships.

Suggested Fields

- Numeric Fields: select or review numeric measure fields to create correlation pairs; avoid ID/index fields because they usually do not explain business relationships

Chart Recommendations

Builds recommended charts from category, date, and value fields and can create dashboards from selected recommendations.

Typical screen area

Category field(s) | Value field(s) | Date field | Search/filter/check boxes | AI | Help | Export

Terms used.

Category = grouping field; Value = numeric measure; Multi-value chart = more than one Y measure; Highest priority = useful unique field combination.

Input entries.

Selections restrict recommendations. If blank, page searches eligible fields. Checkboxes choose charts for dashboard creation.

Model / algorithm functions.

Generates field combinations, filters unsupported chart/value combinations, limits duplicates, assigns priority, validates dashboard-safe chart types, and builds links.

Meaning of output.

Recommended Chart is chart type. Fields show X/category and Y/value choices. Priority ranks usefulness. Add to Dashboard controls dashboard creation.

Sample result.

Chart	Fields	Priority	Dashboard
Column	Region / Sales	Highest	Checked
Line	OrderDate / Revenue	High	Checked

Shortcut.

Use Reset to clear search and selections; Create Dashboard uses current page selections.

Input and Fields Selection

- Category field(s), Value field(s), and Date Field restrict the chart recommendation search; blank selections allow the page to scan possible fields automatically.
- Search filters the recommendation grid, and checkboxes select which rows can be added to a dashboard.

Model and Algorithm

- The page detects category, date, and numeric fields, builds eligible field combinations, scores chart choices, removes weak duplicates, limits the generated list, and marks dashboard-safe charts.
- Highest-priority recommendations balance chart types across valid field combinations.

Output

- The grid shows priority, recommended chart, fields, open chart/open data links, dashboard checkboxes, and a Create Dashboard action for the current page of recommendations.

Why Useful

- Field patterns can be converted into chart suggestions and dashboards.

Suggested Fields

- Category field(s) become chart X/group labels
- Date field supports time charts
- Value field(s) become Y measures

Audit Summaries

Shows which report fields, filters, thresholds, and aggregation options produced each analytical result.

Typical screen area

Selected analytical page options | Report fields | Thresholds/aggregation settings | AI | Help | Export

Terms used.

Audit summary = explanation of result production; Lineage = source fields/settings behind a number; Threshold = cutoff or rule value.

Input entries.

Uses current analytical settings, selected fields, filters, thresholds, aggregation choices, and other options that affect results.

Model / algorithm functions.

Collects selected settings, maps each setting to analytical effect, and builds readable lineage explaining what produced the result.

Meaning of output.

Item identifies setting. Setting/Value show selection. Effect explains how it influenced calculation or filtering.

Sample result.

Item	Setting	Value	Effect
Value Field	Sales	Selected	Used for totals
Threshold	2 Std Dev	Selected	Outlier rule

Shortcut.

Use Audit Summaries when results need to be explained or reviewed.

Input and Fields Selection

- Result Name describes the analytical result being documented.
- Report Fields lists the fields used by the analysis.
- Filters records any filter text, page filter, or business restriction applied.
- Thresholds records numeric limits, readiness thresholds, or rule values.
- Aggregation Options records the aggregation choices used and is locked so the audit remains traceable.

Model and Algorithm

- Audit-summary model documents how an analytical result was produced.
- The page captures selected fields, filters, thresholds, aggregation choices, and record/result counts.
- It turns analysis settings into a repeatable audit row.
- This does not recalculate the business analysis; it records the configuration that produced it.
- The audit can be exported or sent to AI with the visible grid rows.

Output

- Section groups the audit row, such as Report, Analysis, Current Result, or Audit.
- Audit Item names the setting or traceability item being documented.
- Value shows the selected field list, filter, threshold, aggregation option, note, saved result size, user, or timestamp.
- Purpose explains why that audit item matters.
- The grid documents the analytical settings and available session result summaries so the analysis can be reviewed and reproduced.

Why Useful

- Document which fields, filters, thresholds, and aggregation options produced each analytical result.

Suggested Fields

- Audit textboxes: enter the report fields, filters, thresholds, aggregation options, result name, and notes used by the analysis being documented

Map Readiness

Checks whether report data is ready for map use and suggests possible coordinate fields.

Typical screen area

Current report data | Detected latitude/longitude fields | Location/category fields | AI | Help | Export

Terms used.

Latitude/Longitude = coordinates; KML-ready = suitable for map/KML output; Invalid range = latitude outside -90..90 or longitude outside -180..180.

Input entries.

Uses report fields and searches coordinate-like names/values. It ignores index/ID-like fields and suggests latitude/longitude pairs.

Model / algorithm functions.

Detects coordinate candidates, validates missing, duplicate, and out-of-range coordinates, and estimates rows usable for map output.

Meaning of output.

Suggested Fields identify likely coordinates. Records shows affected/ready rows. Status explains good, warning, or review result.

Sample result.

Check	Suggested Fields	Records	Status
Latitude/Longitude	Lat, Long	980	Good
Invalid range	Lat, Long	2	Review

Shortcut.

Run Map Readiness before Map Report to avoid map crashes and bad coordinates.

Input and Fields Selection

- Latitude, Longitude, and optional Name Field dropdowns identify the fields used for map checks.
- The page also scans all fields to suggest likely coordinate columns and excludes ID/index-like fields from coordinate suggestions.

Model and Algorithm

- Map readiness checks missing coordinates, invalid coordinate ranges, duplicate locations, latitude/longitude quality, and KML-ready records.
- Coordinate suggestion scoring looks at field names and usable numeric coordinate ranges.

Output

- The grid shows each readiness check, status, counts, suggested coordinate fields, details, and links to matching records where available.

Why Useful

- Location or coordinate fields can be checked for map and KML readiness.

Suggested Fields

- Latitude field
- Longitude field
- location fields can support geocoding or map labels

Cohort Analysis

Groups records by first activity period and measures later-period activity, retention, record counts, and value.

Typical screen area

Entity field | Date field | Value field | Period | Search | AI | Help | Export

Terms used.

Cohort = entities that started in same period; Retention = percent still active later; Activity period = measured period after first appearance.

Input entries.

Entity identifies customers/users/products/cases. Date determines first and later periods. Value summarizes activity. Period sets month/quarter/year grain.

Model / algorithm functions.

Finds each entity first period, assigns cohort, groups later activity by cohort/activity period, counts entities/records, summarizes value, and calculates retention.

Meaning of output.

Cohort Period is start group. Activity Period is measured period. Period Number is distance from start. Retention % is active/original cohort size.

Sample result.

Cohort Period	Activity Period	Entities	Records	Retention %
2026-01	2026-01	120	310	100%
2026-01	2026-02	84	190	70%

Shortcut.

Use Cohort to study customer, member, product, or case retention over time.

Input and Fields Selection

- Entity field selects the customer, user, account, product, or other entity being followed.
- Date field identifies the activity date used to place records into periods.
- Value field is optional and summarizes value generated by the cohort.
- Period grain controls how dates are grouped, such as month or quarter.
- Search filters source records before cohorts are built.

Model and Algorithm

- Cohort model groups entities by their first observed activity period.
- The first period for each entity becomes the cohort period.
- Later activity is grouped by cohort period and activity period.
- The page counts active entities and records in each period.
- Retention percent compares active entities to the original cohort size.

Output

- Cohort Period shows when the entity first appeared.
- Activity Period shows the later period being measured.
- Period Number shows how far the activity period is from the cohort start.
- Entities, Records, Value, and Retention % summarize cohort behavior.
- Records links open the rows behind each cohort period.

Why Useful

- Customer or user IDs with dates can be grouped into cohorts.

Suggested Fields

- Date Field dropdown: select the first-activity or event date used to assign each entity to a cohort period
- Entity Field dropdown: select customer, account, user, member, order, device, product, or another identifier followed over time

Funnel Analysis

Analyzes ordered statuses or process stages and shows counts, drop-off, conversion percentage, and optional value by stage.

Typical screen area

Stage field | Value field | Stage order | Search | AI | Help | Export

Terms used.

Funnel = ordered process; Stage = one step; Drop-off = records lost since prior step; Conversion % = percent remaining from first step.

Input entries.

Stage contains statuses/steps. Stage order defines sequence when alphabetical order is wrong. Value optionally totals money, quantity, or score.

Model / algorithm functions.

Counts records by stage, applies custom order, calculates drop-off from previous stage, conversion from first stage, and optional value by stage.

Meaning of output.

Step is sequence. Records is count at each stage. Drop Off is loss from previous step. Conversion % shows population remaining.

Sample result.

Step	Stage	Records	Drop Off	Conversion %
1	Lead	500	0	100%
2	Qualified	320	180	64%

Shortcut.

Use Funnel to find bottlenecks in sales, service, approval, or operations.

Input and Fields Selection

- Stage field selects the process step, status, outcome, or funnel stage.
- Optional Value field summarizes revenue, amount, quantity, or other value by stage.
- Stage Order lets the user define the intended process sequence.
- Search filters source records before funnel counts are calculated.
- Use fields such as lead status, order status, conversion step, application stage, or workflow state.

Model and Algorithm

- Funnel model measures movement through ordered stages.
- The page counts records and value for each stage.
- If Stage Order is supplied, that order controls the funnel sequence.
- Drop-off is calculated against the previous stage.
- Conversion percent is calculated against the first stage.

Output

- Step shows the stage order used for the funnel.
- Stage is the selected status or process value.
- Records and Value summarize activity at that stage.
- Drop-off shows how many records were lost from the prior stage.
- Conversion % shows how much of the first stage remains at each step; record links open the stage records.

Why Useful

- Stage/status fields with user/order IDs can show conversion through steps.

Suggested Fields

- Stage Field dropdown: select the status, stage, outcome, step, workflow, or lifecycle field used as funnel steps
- Entity/Record Field dropdown: select customer, order, transaction, account, user, or another identifier counted through stages

ABC Pareto Analysis

Classifies dimensions into A, B, and C groups by cumulative contribution to total value.

Typical screen area

Dimension field | Value field | Aggregate | Search | AI | Help | Export

Terms used.

Pareto = few items drive most value; A class = largest contribution; B = middle; C = long tail; Share = percent of total.

Input entries.

Dimension identifies items such as products/customers/regions. Value is measure to rank. Aggregate controls sum, count, or average behavior.

Model / algorithm functions.

Aggregates value by dimension, sorts descending, calculates share and cumulative share, then classifies A/B/C around 80% and 95% cumulative thresholds.

Meaning of output.

Value is aggregated amount. Share % is item portion of total. Cumulative % adds shares in rank order. ABC Class marks major/middle/tail contributors.

Sample result.

Dimension	Records	Value	Cumulative %	ABC Class
Product A	140	450000	62%	A
Product B	88	135000	81%	B

Shortcut.

Use ABC/Pareto to focus on items driving most of the result.

Input and Fields Selection

- Dimension field selects the category, product, customer, item, or other entity to classify.
- Value field selects the numeric measure used for contribution.
- Aggregation defines how repeated records for the same dimension are combined.
- Search filters records before Pareto ranking is calculated.
- Use sales, revenue, units, profit, cost, or amount fields for meaningful Pareto value.

Model and Algorithm

- ABC/Pareto model finds which dimensions explain most of the total value.
- The page aggregates value by dimension and sorts rows from largest to smallest.
- Share percent is each dimension value divided by total value.
- Cumulative percent is calculated down the sorted list.
- ABC class assigns A to the main contributors, B to the middle contribution, and C to the remaining tail.

Output

- Dimension identifies the category or entity being classified.
- Records links open rows behind that dimension.
- Value, Share %, and Cumulative % show contribution strength.
- ABC Class separates high-impact, mid-impact, and low-impact groups.
- Use this to focus action on the few rows that explain most of the result.

Why Useful

- Find the few categories, products, or customers that explain most of the value.

Suggested Fields

- Category Field dropdown: select product, customer, category, item, department, region, channel, or another dimension to classify by contribution
- Value Field dropdown: select sales, revenue, amount, quantity, profit, cost, or another numeric contribution measure

Data Drift Analysis

Compares field distributions between two selected segments such as periods, groups, locations, or categories.

Typical screen area

Compare field | Segment field | Base value | Compare value | Search | AI | Help | Export

Terms used.

Drift = distribution change; Segment = subset such as period or region; Percentage-point drift = compare share minus base share.

Input entries.

Compare field contains values checked. Segment divides data into groups. Base and Compare values choose the two groups.

Model / algorithm functions.

Counts compare-field values in base and compare segments, converts counts to shares, subtracts base share from compare share, and highlights changed representation.

Meaning of output.

Base/Compare Records are counts. Base/Compare Share % are distributions. Drift Points show percentage-point change.

Sample result.

Field Value	Base Records	Compare Records	Base Share %	Drift Points
Online	210	280	42%	8.5
Retail	190	150	38%	-6.2

Shortcut.

Use Data Drift when report population may have changed between periods or groups.

Input and Fields Selection

- Compare Field selects the value whose distribution is being compared.
- Segment Field selects the field containing base and compare groups or periods.
- Base Value and Compare Value select the two segments to compare.
- Search filters records before drift is calculated.
- Use dates, periods, regions, channels, statuses, or source systems as segment fields.

Model and Algorithm

- Data drift model compares distributions between two selected segments.
- The page counts each Compare Field value in the base segment and compare segment.
- Counts are converted to share percentages inside each segment.
- Drift points are calculated as Compare Share minus Base Share.
- Large positive or negative drift means the distribution changed materially.

Output

- Field Value is the compared category or value.
- Base Records and Compare Records link to the rows on each side.
- Base Share and Compare Share show relative distribution inside each segment.
- Drift Points shows the percentage-point change.
- Use this page to detect data mix changes, channel shifts, product mix changes, or source drift.

Why Useful

- Repeated periods can reveal distribution changes across time.

Suggested Fields

- Date or Segment Field dropdown: select the field that separates base and compare periods or segments for drift review
- Numeric/Value fields: use numeric measures when drift should compare value levels or measure distributions
- Compare Field dropdown: use category/status/product/channel fields when drift should compare distribution changes

KPI Builder

Creates simple calculated business indicators from numeric fields and groups them by a selected dimension.

Typical screen area

Dimension field | Operation | Numerator field | Denominator field | Search | AI | Help | Export

Terms used.

KPI = key performance indicator; Numerator = top value in ratio; Denominator = bottom value; Ratio = numerator divided by denominator.

Input entries.

Dimension groups the KPI. Numerator and denominator supply numeric totals. Operation chooses ratio, difference, sum, or product. Search filters rows.

Model / algorithm functions.

Sums numerator and denominator by dimension, applies selected operation, handles divide-by-zero safely, and builds KPI value for each group.

Meaning of output.

Numerator/Denominator show aggregated inputs. Operation explains formula. KPI Value is calculated indicator. Records link to rows behind group.

Sample result.

Dimension	Records	Numerator	Denominator	KPI Value
West	120	240000	8000	30
East	95	171000	5700	30

Shortcut.

Use KPI Builder for revenue per unit, cost per case, margin rate, or productivity.

Input and Fields Selection

- Dimension field groups KPI results by category, department, customer, product, period, or location.
- Numerator field is the main numeric measure.
- Denominator field is used for ratios or comparisons when selected.
- Operation chooses ratio, difference, sum, product, or other KPI calculation.
- Search filters the source records before KPI values are calculated.

Model and Algorithm

- KPI model builds reusable business metrics from selected fields.
- The page groups records by dimension.
- Numerator and denominator values are summed for each group.
- The selected operation is applied to produce the KPI value.
- Division-style calculations protect against missing or zero denominator values.

Output

- Dimension identifies the KPI group.
- Records links open rows behind the KPI.
- Numerator and Denominator show the summarized inputs.
- Operation explains how KPI Value was calculated.
- KPI Value is the resulting metric for monitoring performance or comparing groups.

Why Useful

- Numeric measures can become KPIs, totals, averages, rates, and thresholds.

Suggested Fields

- Numerator / KPI value fields: select the main numeric measure used to build the KPI
- Dimension Field dropdown: optionally group KPI results by category, product, customer, department, location, or period
- Date Field dropdown: optionally select a date field when KPI results should be reviewed by period

Data Dictionary

Generates field-level documentation with detected type, examples, blanks, distinct values, summaries, and recommended analytical use.

Typical screen area

Field group | Examples | Search | AI | Help | Export

Terms used.

Data dictionary = field documentation; Detected type = inferred numeric/date/text behavior; Recommended use = suggested role such as category, date, measure, or ID.

Input entries.

Field group filters all/numeric/date/text fields. Examples controls sample values. Search narrows fields by name.

Model / algorithm functions.

Scans fields, detects likely type, counts blanks/distinct values, collects examples, calculates numeric statistics, and suggests analytical roles.

Meaning of output.

Field identifies column. Detected Type explains usage. Blanks/Distinct describe completeness and uniqueness. Recommended Use helps choose fields.

Sample result.

Field	Detected Type	Blanks	Distinct Values	Recommended Use
Sales	Numeric	0	923	Measure / KPI
OrderDate	Date	2	356	Time grouping

Shortcut.

Use Data Dictionary before sharing reports or sending context to AI.

Input and Fields Selection

- Field Group filters the dictionary to all fields, numeric fields, date fields, text fields, or other supported groups.
- Number of Examples controls how many sample values are shown.
- Search filters by field name or related field text.
- The page uses all available current report columns.
- Use this page before deeper analysis when field meaning or data behavior is unclear.

Model and Algorithm

- Data dictionary model documents field meaning and practical analytical use.
- The page scans values to detect type, blanks, distinct values, examples, and numeric summaries.
- It identifies likely measures, categories, IDs, dates, and text fields.
- Recommended use is inferred from field behavior and naming patterns.
- The dictionary supports choosing better fields on other analytics pages.

Output

- Field and Detected Type identify each column.
- Records, Blanks, and Distinct Values describe completeness and uniqueness.
- Min, Max, Average, and Std Dev appear where applicable.
- Examples show representative values from the data.
- Recommended Use explains how the field can be used in analytics, charts, filters, grouping, or quality review.

Why Useful

- Field-level documentation is useful for any unfamiliar dataset.

Suggested Fields

- Field Group / Search: use all fields, or narrow to fields whose meaning, type, examples, blanks, or distinct values need documentation

Data Readiness Scanner

Scans the current report or imported dataset and recommends useful analytics, market models, charts, maps, and quality checks by readiness score.

Typical screen area

Search | Build | Reset | AI | CSV | Excel | Recommended Analytics Grid | Open | Records

Terms used.

Readiness score = algorithm-assigned fit of an analysis to detected fields; Suggested Fields = guidance for the dropdowns and input controls on the recommended page.

Input entries.

Uses the current report result or imported dataset in memory. No manual field selection is required. Search optionally narrows recommended analysis rows.

Model / algorithm functions.

Detects numeric, date, category, ID, customer, order, product, location, price, quantity, revenue, and status fields; counts missing values and duplicate records; then assigns readiness scores.

Meaning of output.

Analysis identifies the suggested page. Readiness and Score show relative fit. Why Useful explains the recommendation. Suggested Fields tells what fields make sense in that page's dropdowns and controls.

Sample result.

Analysis	Readiness	Score	Suggested Fields	Open
Detail Analytics	High	95	Groups: Region, Product; Value: Sales	open
Overall Statistics	High	95	Numeric: Sales; Text: Region	open
Fields Correlation	High	90	Numeric pairs: Sales, Profit	open
Market Pricing	Possible	82	Price: UnitPrice; Response: Quantity	open

Shortcut.

Open Data Readiness Scanner first when a dataset is unfamiliar, then begin with High readiness recommendations.

Input and Fields Selection

- Uses the current report or imported dataset as an unknown table.
- No manual field selection is required; the scanner evaluates every column.
- Search can narrow the visible recommendation grid after scoring.
- The scanner detects numeric measures, dates, categories, IDs, products, customers, orders, locations, prices, quantities, revenue, and status/outcome fields.
- Detected fields are used to suggest which analytics and market models are likely useful.

Model and Algorithm

- Readiness model scores each analysis by checking whether the dataset contains the fields normally required.
- Column names, data types, blank counts, duplicate records, distinct values, and field combinations are inspected.
- High readiness means the minimum field pattern exists and the analysis should likely produce useful output.
- What Next links are generated from known relationships between pages.
- The Suggested Work Flow connects high-readiness pages into practical analysis paths.

Output

- Readiness Level and Score rank the usefulness of each page for the current dataset.
- Why Useful explains the reason the analysis is recommended.
- Suggested Fields names the dropdowns or input controls on that page and the fields that make sense to use.
- Open links launch the recommended page; What Next links show follow-up analysis.
- Records links return to Data Explorer; the workflow line suggests an efficient order for reviewing high-scored pages.

Why Useful

- Tells users which analysis pages make sense and why each page is recommended for the current data.

Suggested Fields

- Suggested Fields names the dropdowns or input controls on each page and the fields that make sense to use.