[DISTRIBUTIONAL]

Distributional's Analytics Workflow

Intro

Distributional has updated and continues to update our core functionality.

See docs.dbnl.com for the latest on our data pipeline, analytics workflow, and enterprise platform.

Production AI behavior is a black box, which leaves AI teams struggling to improve, fix, and scale their products. Distributional's adaptive analytics platform empowers these teams to have confidence in AI behavior—the interplay and correlations between users, context, tools, models, and metrics. The product (DBNL) provides a detailed snapshot of aggregate AI product behavior, regularly surfaces behavioral signals, links evidence to facilitate deep root cause analysis, and enables teams to track relevant signals over time. This empowers AI teams to better understand the behavior of their users and AI products so that they can fix and improve those products with confidence.

Distributional's adaptive analytics product includes an analytics workflow that is supported by a data pipeline and powered by an enterprise platform. The focus of this paper is the analytics workflow that empowers AI product teams to continuously discover, investigate, and track insights on their AI products and agents so they can fix and improve them over time.

Table Of Contents

Workflow

Distributional's analytics workflow empowers AI product teams to discover, investigate and track insights on their AI products and agents. The more you use this workflow, the more tailored it becomes to your AI product, serving as a way to align the team on prioritization of decisions to improve or fix the product as it scales.

If you are a team running AI in production at scale that is struggling to understand usage patterns and needs systematic data analysis to make better product decisions, DBNL is designed for your circumstances. This workflow makes it possible to uncover entirely new insights on interesting usage patterns, outliers, changes, or other behavioral signals that can drive continuous improvement of your AI product or agent.

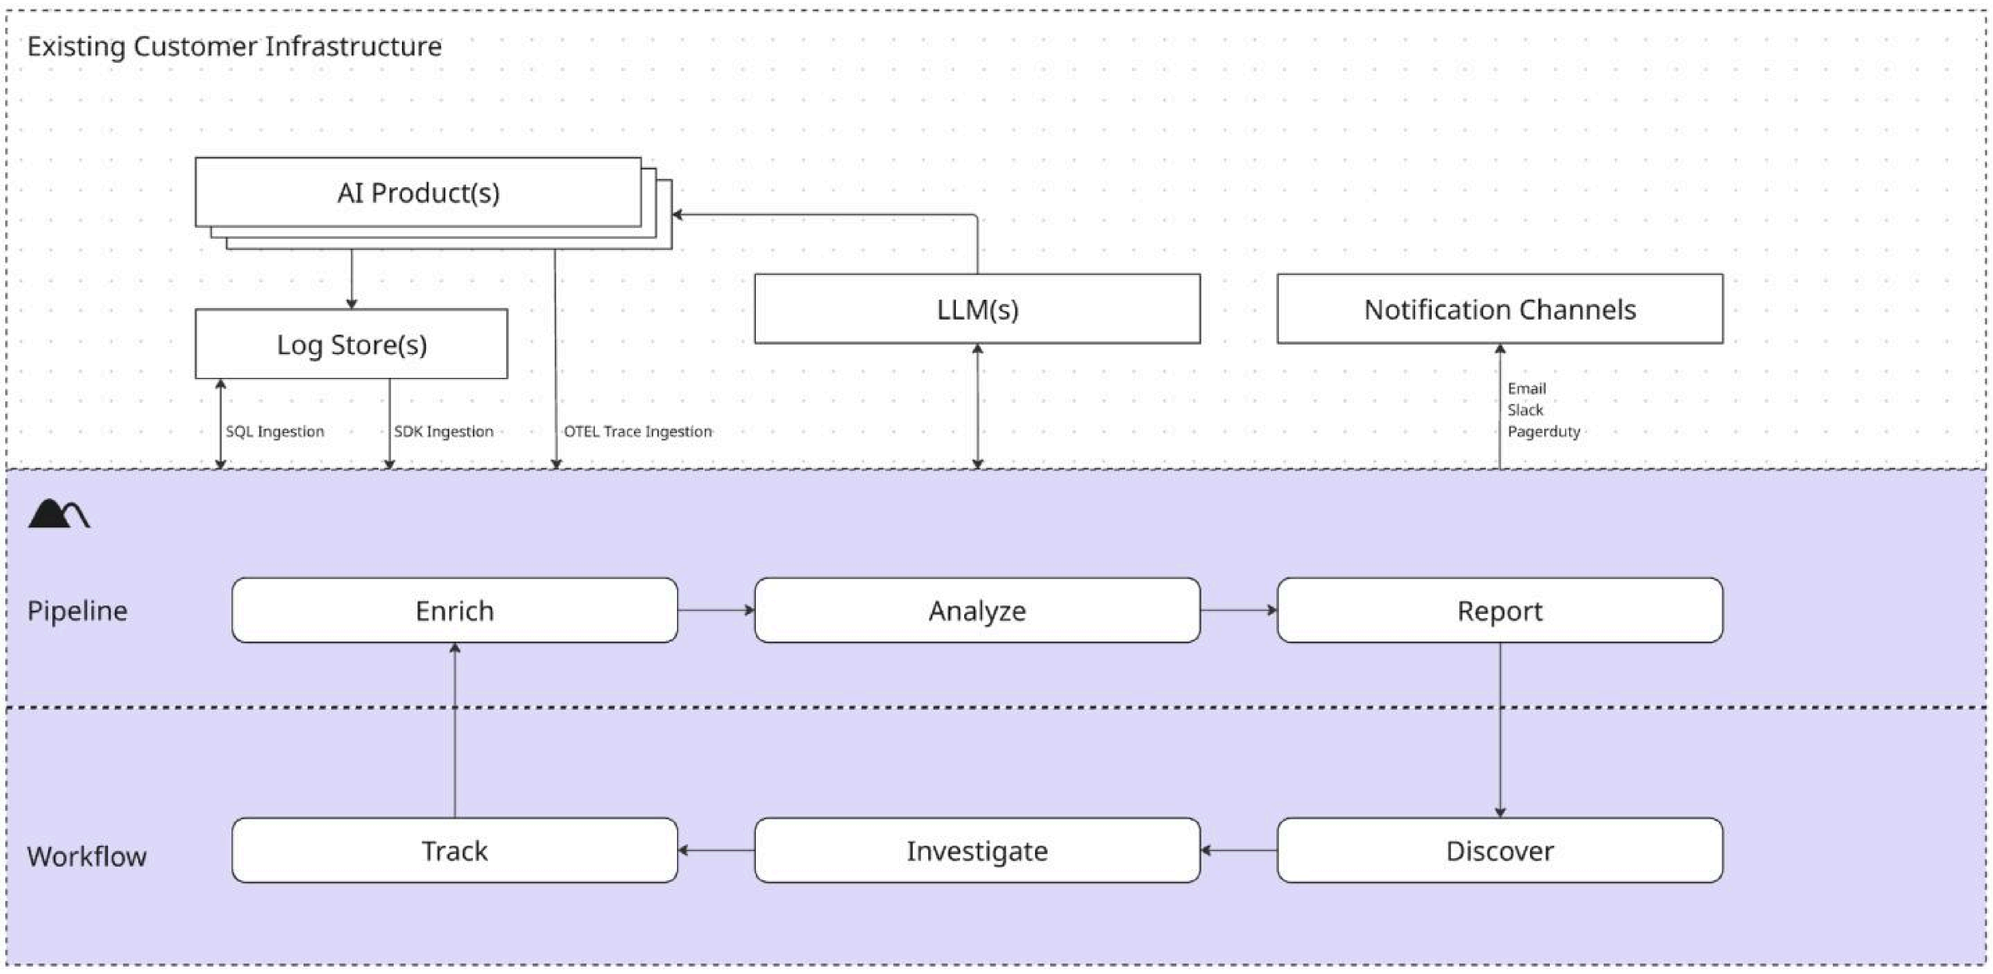

Discover Insights On Production AI

Distributional's analytics workflow starts with daily insights on your AI products and agents. The Distributional Data Pipeline enriches your production AI Logs, analyzes these enriched Logs to produce behavioral signals, and publishes these signals as human readable Insights in your dashboards and notifications channels. These Insights are based on outliers, shifts, changes, topics, clusters, or other interesting usage patterns that represent opportunities to improve or fix your AI product.

INSIGHTS

An Insight is a human-readable explanation and quantification of behavioral signals generated from unsupervised analysis of enriched Logs as part of the Distributional Data Pipeline. Distributional's pipeline produces Insights when filters on Logs correspond to unique behavior. Insights represent relevant signals that the user can triage and refine through the Explorer or inspection of Logs and track as Metrics or Segments.

Most projects generate 5-20 new Insights per week. Projects with stable, consistent behavior may generate fewer Insights, while projects with volatile or rapidly changing behavior may generate more. Distributional requires at least 7 days of data to establish behavioral baselines and produce Insights related to topics.

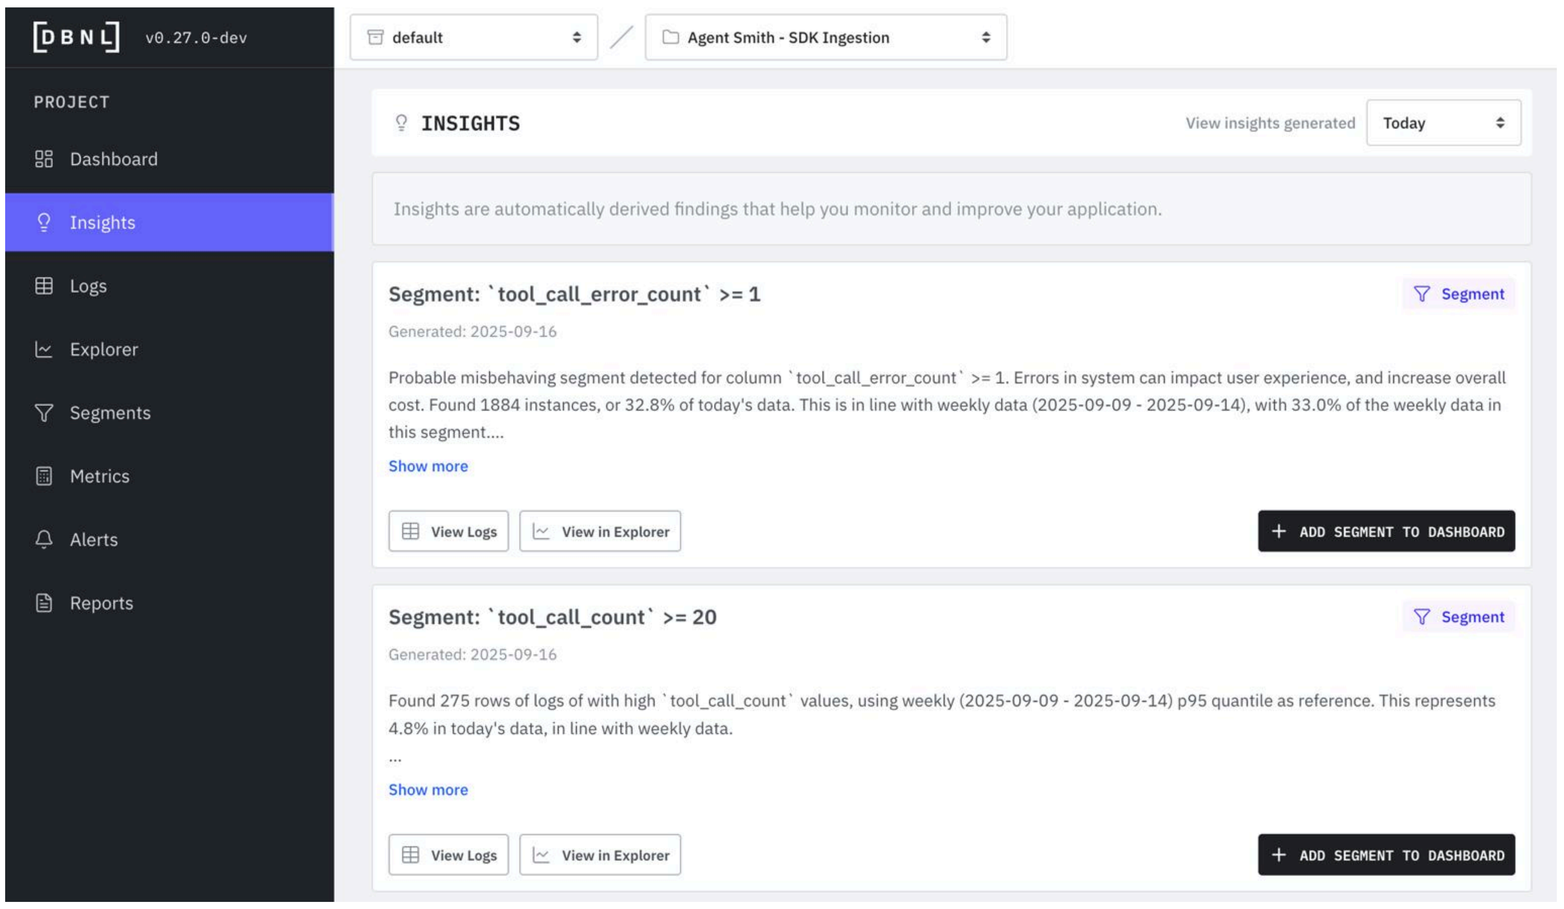

DASHBOARDS

Insights are published to Distributional Dashboards. Dashboards are collections of histograms, time series, and statistics of Columns, Segments, and Metrics. These Dashboards serve as a resource to align AI product teams on how their AI products are being used and identify ways to improve or fix these products over time.

There are a few default Dashboards for each Project:

- Monitoring Dashboard: Recommended graphs and statistics built from Columns and data from the DBNL semantic convention
- Agent Dashboard: Directed graph of all agent tool calls as a Sankey chart to allow for rapid exploration of tool call chains
- Segments Dashboard: Count graphs and statistics for all tracked Segments
- Metrics Dashboard: Histograms, time series, and statistics of generated metrics

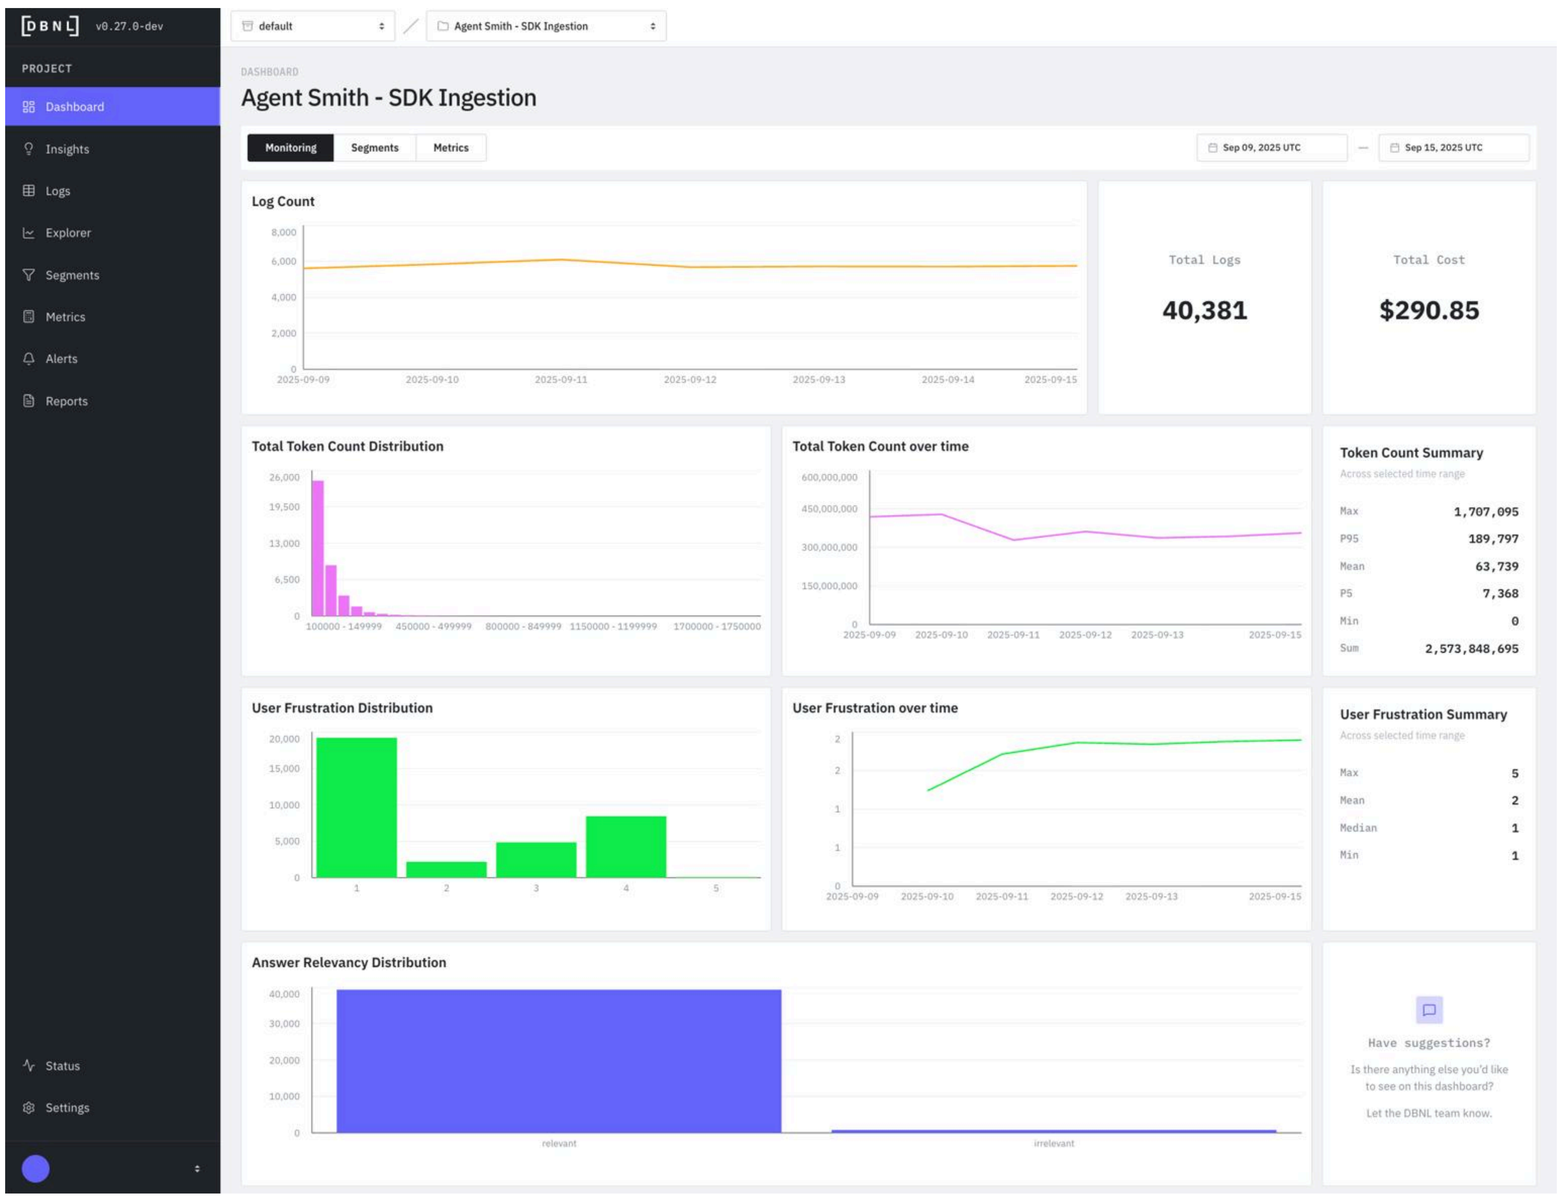

NOTIFICATIONS

Insights are also published to Distributional's Notification Connections. These Notification Connections empower you to connect your preferred notification tools so you are informed when there are new behavioral signals or meaningful changes to expected usage patterns. Distributional supports Slack and email with more channels coming soon.

The goal of these Notification Channels is to ensure you stay informed on your AI products in the places where you are already collaborating or receiving regular product updates.

Investigate These Insights With Evidence

Each new Insight is linked to supporting evidence to facilitate rapid root cause analysis. This evidence includes a curated set of Logs that are most correlated with the insight. It also includes single Segment analysis, Segment comparisons, and temporal comparisons in the Explorer.

LOGS

Distributional's Logs are filterable subsets of all ingested data and all generated metrics. The Logs page allows the user to inspect specific Logs with certain properties defined by a specific time window and specific filters on Columns or Metrics. Typically, the Logs page is visited as part of investigating a specific Insight or by clicking on part of a chart from a Dashboard, in which case the filters and time window will already be applied.

The Logs page allows the user to inspect specific Logs with certain properties defined by

- A specific time window (default: last 7 full days of data)
- Specific filters on Columns or Metrics (default: no filters)

Typically, the Logs page is visited as part of investigating a specific Insight or by clicking on part of a chart from a Dashboard, in which case the filters and time window will already be applied.

Individual Logs can be viewed in a variety of ways:

- Detailed View: All Columns and Metrics of the log viewed together and optionally expanded.
- Trace View (if spans provided): The waterfall trace view of latency and timing for each individual span.
- Session View (if session_id provided): All associated Logs for the given session, along with Metrics.

As part of inspecting the Logs the user can view the filtered Logs as charts and tables in the Explorer and save the specific filters as a Segment to publish it on the Segments Dashboard.

EXPLORER

The Explorer enables rapid analysis and triage of Segments by performing graphical and statistical comparison between different subsets of Logs over different time windows and/or filters. As new data is published to dashboards and channels, Explorer updates as well.

There are three main types of exploration afforded by the Explorer:

- Single Segment: Quickly see all Metrics for a given time window and single filter on the Logs. This allows for an aggregate view of all Metrics.
- Segment Comparison: Compare two different filters on the Logs or a filter and its complement across the same time window. This allows for comparison of Metrics between filters or between filters and the rest of the Log data.
- Temporal Comparison: Compare a single filter across two adjacent time windows. This allows for a Metric comparison of before/after for a given Segment.

When investigating an issue, start with the time series to identify when it started, then use the histogram to understand what values are problematic, and finally check the Logs page to see which specific Logs exhibit the behavior.

Track Relevant Signals Over Time

The purpose of these Insights is to guide prioritization of fixes and improvements to your AI product or agent. As you investigate these Insights and decide which are most relevant, you can track these Insights as Metrics or Segments. When tracked, these Metrics and Segments are automatically analyzed in the Distributional Data Pipeline and become a part of the Distributional Dashboards going forward.

METRICS

Standard Metrics are functions that can be computed using non-LLM methods. They can be built using the functions available in the DBNL query language. Common examples are response length, input complexity, question mark detection, or string similarity.

LLM-as-judge Metrics call an LLM to assess production AI Logs against a specified attribute. Each LLM-as-judge metric is one of two types:

- Classifier Metric: Outputs a categorical value equal to one of a predefined set of classes.
- Scorer Metric: Outputs an integer in the range [1, 2, 3, 4, 5].

Use standard Metrics when you need fast, deterministic calculations (e.g., word counts, text length, keyword matching, readability scores). Use LLM-as-judge Metrics when you need semantic understanding (e.g., relevance, tone, quality, groundedness).

As you add Metrics over time, Distributional's Data Pipeline automatically computes them daily from production AI Logs.

SEGMENTS

Segments are saved filters on Logs corresponding to a specific behavioral signal discovered manually or from an Insight.

You create a Segment when you've identified a meaningful behavioral pattern you want to track over time, such as:

- Error conditions: Logs containing specific error types or failure patterns
- High-value interactions: User sessions with purchases, conversions, or key actions
- Quality issues: Low-scoring responses that need monitoring
- User cohorts: Specific user groups (power users, new users, etc.)
- Performance bottlenecks: Requests exceeding latency thresholds

Once saved, Segments are automatically analyzed in future pipeline runs, generating dedicated Metrics and appearing on Dashboards.

Adapt These Analytics To Your AI Products

As you use DBNL over time, these Segments, Metrics, and Insights become tailored to your specific AI product or agent, improving the analysis and increasing the capacity for your team to make and track product decisions. In this sense, the workflow, as well as the analytics supporting it, adapt to your preferences over time.

With Distributional's adaptive analytics platform, customers can complete the AI product feedback cycle. Empowered by continuous Insights from DBNL, teams can understand usage patterns of their AI products, and use this information to improve and fix their products over time.

When investigating an issue, start with the time series to identify when it started, then use the histogram to understand what values are problematic, and finally check the Logs page to see which specific Logs exhibit the behavior.

Learn more

- docs.dbnl.com
- distributional.com/blog

[DISTRIBUTIONAL]

About Distributional

Distributional's platform is designed to give customers complete control and optionality over their data and integrations. The platform can adapt to specific security, privacy, and regulatory needs, while still providing the quality and insights necessary to test AI applications.

If you'd like to learn more, reach out to the Distributional team. We're here to help!

Learn more at distributional.com

Follow us on:

- LinkedIn
- YouTube