

CATEGORY	FEATURE	WHAT IS IT?	WHY IS IT INCLUDED?	HOW DOES IT WORK?
Unstructured Source Integration	Public Data Integration	<ul style="list-style-type: none"> <li>Processes unstructured text from news, regulatory, corporate and other public web sources</li> </ul>	<ul style="list-style-type: none"> <li>Turns public unstructured data sources into a manageable, actionable data resource for client use cases and business processes</li> </ul>	<ul style="list-style-type: none"> <li>Integrates directly with RSS feeds</li> <li>Supports imported documents in PDF, HTML and more source formats</li> <li>Uses a cloud-hosted schema-less NoSQL data lake to scale and store immense volumes of unstructured data</li> </ul>
	Private Data Integration	<ul style="list-style-type: none"> <li>Allows a direct, secure import of private documents for analysis</li> </ul>	<ul style="list-style-type: none"> <li>Turns your private unstructured data sources into a manageable, actionable data resource for client use cases and business processes</li> <li>Supports the data privacy expectations of enterprise-scale organizations</li> </ul>	<ul style="list-style-type: none"> <li>Uses a private cloud-hosted schema-less NoSQL data lake to scale and store immense volumes of unstructured data</li> </ul>
	Document Preparation	<ul style="list-style-type: none"> <li>Allows for the import of scanned and other non-machine readable documents</li> <li>Detects the language used throughout a document</li> </ul>	<ul style="list-style-type: none"> <li>Prepares documents for additional downstream processing activities required for use cases and business processes</li> </ul>	<ul style="list-style-type: none"> <li>Uses Optical Character Recognition (OCR) capabilities to prepare scanned documents for further processing</li> <li>Leverages computer vision capabilities to isolate tables for specific processing activities</li> </ul>
Document Classification	Document Classification	<ul style="list-style-type: none"> <li>Interprets individual documents to provide a configurable classification</li> <li>Analyzes a group of documents to identify anomalies or outliers</li> </ul>	<ul style="list-style-type: none"> <li>Classifies documents to support a variety of business use cases and downstream processing requirements</li> <li>Classification and grouping of documents more efficiently supports document digitization programs and enterprise document search requirements</li> </ul>	<ul style="list-style-type: none"> <li>Utilizes an array of techniques such as Word2vec, Doc2vec, TF-IDF, and Latent Dirichlet Allocation (LDA) to support document classification requirements</li> </ul>
Language Modeling & Training	Specialized Language Modeling	<ul style="list-style-type: none"> <li>Rapidly train new models to fit your source, client, or use case</li> <li>Trained models achieve 90%+ accuracy across a variety of news, financial and corporate documents, including legalese-filled case documents and legal filings</li> </ul>	<ul style="list-style-type: none"> <li>Different models can be applied to different types of document sources, supporting unique, fine-tuned analysis across legal documents, financial documents, news articles and more</li> </ul>	<ul style="list-style-type: none"> <li>Uses general and specialized neural-language models to better support disparate sources of information</li> </ul>
	Model Feedback & Tuning	<ul style="list-style-type: none"> <li>Provide feedback to the model for continued model improvement in support of higher accuracy requirements</li> </ul>	<ul style="list-style-type: none"> <li>Allows user influence for language model tuning to enable continuously improving accuracy of data identification and extraction across source documents</li> </ul>	<ul style="list-style-type: none"> <li>Leverages machine learning capabilities to continuously improve language models for greater accuracy when processing the language of specific source documents</li> </ul>

CATEGORY	FEATURE	WHAT IS IT?	WHY IS IT INCLUDED?	HOW DOES IT WORK?
Data Identification & Extraction	Entities & People	<ul style="list-style-type: none"> <li>• Uses trained language models to identify, tag and extract entities and people within strings of texts, tables and other semi-structured text</li> <li>• Pre-existing fully-trained models identify entities and people throughout most sources and language patterns</li> <li>• Customizable models can be added, configured and trained in a matter of days or weeks to support unique use cases</li> </ul> <p>Supports a variety of customizable business rules to extract and prepare data for downstream processing</p>	<ul style="list-style-type: none"> <li>• Identifies entities and people within unstructured text for exploration within the user interface</li> <li>• Extracts entities and people from unstructured text for integration with downstream data activities and business processes</li> </ul>	<ul style="list-style-type: none"> <li>• Executes custom identification that works with parts of speech, lexemes and syntactic parse</li> <li>• Understands context to identify pronouns, family names, entity names, abbreviations and more</li> </ul>
	Events	<ul style="list-style-type: none"> <li>• Uses a mixture of trained and configurable language models to identify, tag and extract events within strings of texts, tables and other semi-structured text</li> <li>• Pre-existing models identify a variety of corporate action and risk related events within a document, such as name changes, merger/acquisitions, director/shareholder changes, penalty/fines, expulsions/suspensions, insolvency and more.</li> <li>• Models for new events can be added, configured and trained in a matter of days or weeks to support unique use cases</li> <li>• Supports a variety of customizable business rules to extract and prepare data for downstream processing</li> </ul>	<ul style="list-style-type: none"> <li>• Events provide valuable insights specific to your use cases and business processes</li> <li>• Triangulate events, entities and people to develop context within your unstructured source documents</li> </ul>	<ul style="list-style-type: none"> <li>• Executes custom identification that works with parts of speech, lexemes and syntactic parse</li> <li>• Understands context to identify events that are attributable back to specific entities, people and data attributes</li> </ul>

CATEGORY	FEATURE	WHAT IS IT?	WHY IS IT INCLUDED?	HOW DOES IT WORK?
Data Identification & Extraction	Data Attributes	<ul style="list-style-type: none"> <li>• Uses a mixture of trained and configurable language models to identify, tag and extract additional data attributes within strings of texts, tables and other semi-structured text</li> <li>• Pre-existing models identify a variety of popular data attributes within a document, such as dates, locations, currency, time and more</li> <li>• Models for new data attributes can be added, configured and trained in a matter of weeks (longer for complex data attributes)</li> <li>• Supports a variety of customizable business rules to extract and prepare data for downstream processing</li> </ul>	<ul style="list-style-type: none"> <li>• Supports additional use cases and business processes related to specific, unique data attributes extracted from unstructured source documents</li> <li>• Associate data attributes to entities, people and events to develop context within your unstructured source documents</li> </ul>	<ul style="list-style-type: none"> <li>• Executes custom identification that works with parts of speech, lexemes and syntactic parse</li> <li>• Understands context to identify data attributes that are attributable back to specific entities, people and events</li> </ul>
	Tabular Data	<ul style="list-style-type: none"> <li>• Identifies and extracts data nested within tables in unstructured source documents</li> <li>• Pre-existing table processing models identify entities, people and events, as well as a variety of popular data attributes within tables</li> <li>• Models for custom tables with specific accuracy requirements can be configured and trained in a matter of weeks (longer for complex tables)</li> <li>• Supports a variety of customizable business rules to extract and prepare table data for downstream processing</li> </ul>	<ul style="list-style-type: none"> <li>• Supports data identification and extraction beyond natural language, deriving value out of the content within tables in an unstructured source document</li> <li>• Associate table data to entities, people and events to develop context within your source documents</li> </ul>	<ul style="list-style-type: none"> <li>• Executes an ensemble of algorithms and heuristics to provide context while extracting entities, people, events and data attributes from tables within unstructured source documents</li> <li>• Understands context to identify table data that is attributable back to specific entities, people and events</li> </ul>

CATEGORY	FEATURE	WHAT IS IT?	WHY IS IT INCLUDED?	HOW DOES IT WORK?
Text Analytics UI	Document SpeedRead	<ul style="list-style-type: none"> <li>Summarizes all relevant entities, people, events and data attributes within the text, enabling faster access to specific sections, words or headlines of a document</li> </ul>	<ul style="list-style-type: none"> <li>Allows efficient document exploration within a secure web application</li> <li>Improves the speed of reading and comprehension to streamline manual data operational activities for a variety of client use cases and business processes</li> </ul>	<ul style="list-style-type: none"> <li>Web application UI is surfaced securely for browser-based access by users</li> <li>Intuitive and modern UI enables efficient usability with regards to manual data operations tasks</li> </ul>
	Explore & Discover UI	<ul style="list-style-type: none"> <li>Explores and searches all unstructured source content within the data lake, filtering content and documents within a specific time range and for focused organizations, people and events</li> </ul>	<ul style="list-style-type: none"> <li>Provides a real-time, manageable overview of all unstructured data sources within the data lake</li> <li>Allows users to explore the data lake content, narrow document scope and improve operational efficiency for a variety of client use cases and business processes</li> </ul>	<ul style="list-style-type: none"> <li>Schema-less data lake supports complex search and filtering in near real-time</li> <li>Web application UI is surfaced securely for browser-based access by users, enabling efficient usability for manual data operations tasks</li> </ul>
	Advanced Analytics Dashboard	<ul style="list-style-type: none"> <li>Summarizes all unstructured content within the data lake in a variety of dashboards and UI widgets to provide valuable insight into the trends of organizations, people and events mentioned within content</li> </ul>	<ul style="list-style-type: none"> <li>Provides analytical meta-data useful to administrative users</li> <li>Exposes trends of data mentioned within unstructured content in support of specific business processes and use cases</li> </ul>	<ul style="list-style-type: none"> <li>Uses widget-based charts and dashboards driven by scalable APIs</li> <li>Web application UI is surfaced securely for browser-based access by users, enabling efficient usability for monitoring data trends within the unstructured source content processing</li> </ul>
	Administration & Configuration UI	<ul style="list-style-type: none"> <li>Enables UI-based configuration of key sources and events, enabling flexibility for a variety of client use cases and business processes</li> <li>New unstructured data sources and events can be onboarded and tailored for use cases in a matter of days</li> </ul>	<ul style="list-style-type: none"> <li>Enables new unstructured data sources to be onboarded quickly for evolving business processes and use cases</li> <li>Exposes a human-readable model for configuring new events, which are processed across source content to highlight actionable insights</li> </ul>	<ul style="list-style-type: none"> <li>Web application UI is surfaced securely for browser-based access by administrators, enabling efficient onboarding of sources and events for specific use cases across unstructured content</li> <li>Configuration information for sources and events is leveraged by the back-end Text Analytics engine to process new unstructured content for insights important to client use cases</li> </ul>



CATEGORY	FEATURE	WHAT IS IT?	WHY IS IT INCLUDED?	HOW DOES IT WORK?
Exports, Alerts & Integrations	Data Delivery Module	<ul style="list-style-type: none"> <li>Subscribes the user to user-selected entities, people and/or events processed within unstructured source content in the data lake</li> <li>Supports system-to-system integration</li> </ul>	<ul style="list-style-type: none"> <li>Refines and tailors the processed content to just the entities, people, events and data attributes important to your use cases</li> <li>Enables downstream system integrations for improved business processes</li> </ul>	<ul style="list-style-type: none"> <li>Tailored, use case-specific entity and people lists can be integrated through back-end integrations for refining the processed content</li> <li>Downstream consumption of content and data is available through batch, REST APIs and message queues</li> </ul>
	Notifications & Alerts	<ul style="list-style-type: none"> <li>Provides notifications when subscribed entities, people and/or events are detected within processed source content</li> </ul>	<ul style="list-style-type: none"> <li>Ensures users and downstream systems are immediately aware of events and insights processed within source content for the entities and people that matter most to your business</li> </ul>	<ul style="list-style-type: none"> <li>Supports notifications directly within the application or in downstream systems via message queuing architecture</li> <li>Email notifications can be scheduled or supported dynamically</li> </ul>
Cloud & DevSecOps	Cloud & DevSecOps	<ul style="list-style-type: none"> <li>Scalability, security, flexibility and a mix of best practices including continuous integration and continuous delivery (CI/CD) and an API-driven architecture optimized for cloud</li> </ul>	<ul style="list-style-type: none"> <li>Backs the overall Text Analytics SaaS solution with the security and development best practices that enterprise-scale organizations expect of their software solutions</li> <li>Cloud optimization enables the solution to scale to a near infinite amount of source content and processing for the most demanding of use cases</li> </ul>	<ul style="list-style-type: none"> <li>Best-of-breed DevSecOps technologies are brought together to automate the build/secure/test/deploy development lifecycle</li> <li>Cloud optimization in AWS ensures scalability and security needed for even the most complex and demanding use cases and performance requirements</li> </ul>