



LUMINOSO

The buyer's guide to text analytics

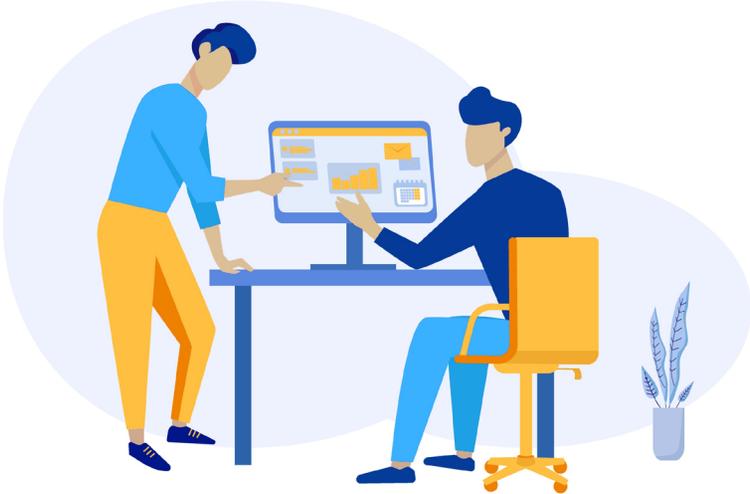
What to ask before you buy

A Luminoso eBook



Table of contents

- Introduction 3
- Four questions to ask before buying 4
 - 1. How much time do I need for setup and maintenance? 4
 - 2. What kind of data do I need, and how much of it? 5
 - 3. How powerful and sophisticated are the insights I'll get? 6
 - 4. What features does this product offer? 7
- Ask, evaluate, choose. Then, get started. 9
- The text analytics buyer checklist 10



Introduction

If you're involved with your organization's customer or employee experience initiatives, you're familiar with the value of the feedback you're receiving – and the pain of analyzing it. No matter how well you design a structured set of questions, the best insights often come from responses to open-ended questions.

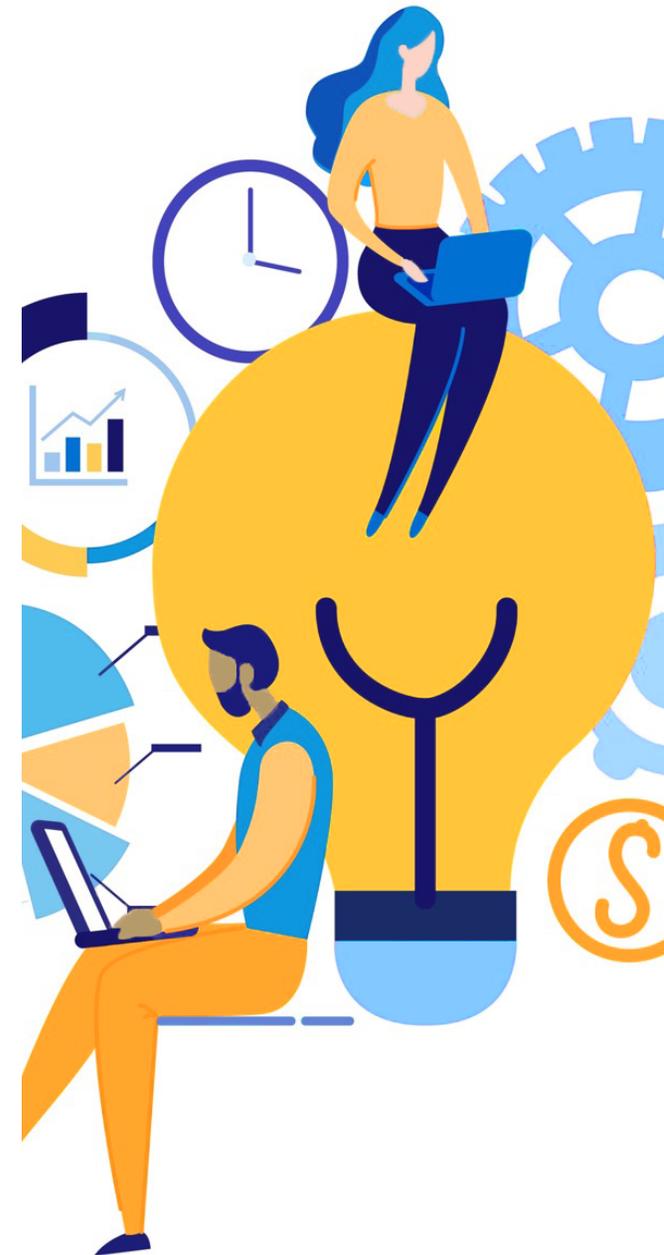
Natural language understanding (NLU) combines the power of artificial intelligence (AI), computational linguistics, and computer science to allow a machine to understand natural language in a “human” context. Text analytics apply NLU to read text more quickly and accurately than humans can. Common features of text analytics tools include sentiment analysis, summarization, and classification – a supposedly perfect solution to your open-ended text challenge.

While text analytics methods continue to become increasingly effective, most widely-available solutions still require setup and tuning to guarantee results, especially in industries with a lot of specific or rapidly changing vocabulary. Unfortunately, open-ended text commonly contains lots of organization or industry-specific jargon, slang, misspellings, and complex language. And although many text analytics tools handle expected words well, these complexities – which hold the most valuable insights – can cause them to stumble.

Inevitably, while handling open-ended questions, you've realized that there must be some way to automate this analysis and avoid digging through data manually.

And you're correct: there are text analytics solutions out there that can deliver on these promises, in terms of both usability and power.

This guide provides a framework for you to evaluate current text analytics options and ensure you're making the best investment for your organization.



Four questions to ask before buying

So you've decided to stop manually analyzing the text you have laying around. Or, you have a tool in place you're not quite happy with, and you've committed to making a change.

Realizing there's an easier, more effective way to analyze feedback is the first step in your text analytics journey. Now, let's get you on the path that's best for your organization by asking the right questions.

1. How much time do I need for setup and maintenance?

Here comes a difficult truth: most commercially-available text analytics solutions require work before the system works for you. This means anything from building ontologies and libraries, to keyword tagging and data training ... and even coding, manual updates, or maintenance. If your team or organization is small, or strapped for developers and data scientists, you may not have hours to perfect a text analytics tool. And even if you have the in-house resources, your data and development experts shouldn't be relegated to solution setup and maintenance.

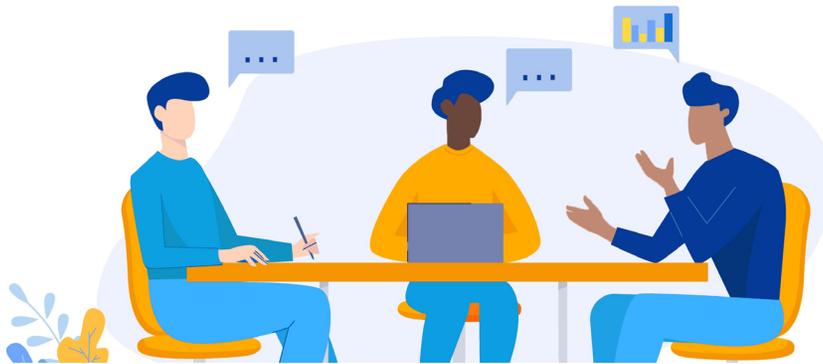
Don't be discouraged. No-code text analytics are here. This means you and your team can get up to speed and share findings quickly, without technical experience. No-code systems do not require developers before machine learning can happen.

Factors to consider:

- The amount of manual work the solution requires before machine learning can take place
- Whether or not your organization needs a dedicated data scientist to use the system
- How much specialized code the solution requires
- If business users need technical training to easily navigate the application

2. What kind of data do I need, and how much of it?

What your organization does and how you analyze and report on data determines a lot about your approach to buying a solution. You could have hundreds or millions of documents to analyze on a cadence that could vary from daily to yearly. The savvy text analytics buyer has an intimate knowledge of the balance between volume and analysis frequency.



Datasets that contain fewer than 1000 documents are too small for many solutions. If your dataset is this size, choose a tool that doesn't need thousands of examples to train it – or any training at all. And even if you're examining over 1000 documents with each upload, if your uploads aren't frequent, make sure that you're not overpaying for a solution that's meant for a buyer with a high frequency of analysis.

If your volume is high and your analysis frequency is often, look for a solution that lets you home in on deep insights, fast.

Factors to consider:

- The number of documents needed to train the system before it can understand domain-specific language
- Whether the solution's pricing is volume-based, and if so, what tiers are available
- The time and effort it takes go from a raw data upload to meaningful insights

3. How powerful and sophisticated are the insights I'll get?

Once you've settled on a solution that meets your basic needs, it's time to ask tougher questions. Even if a product doesn't require a ton of up-front work to get started, you also need to avoid reaching a limit on how deep you can dig in your data. Many tools can't handle domain- or industry-specific words or new and emergent terms they haven't seen, creating blind spots in language analysis. Ensuring that static, ontology-based solutions don't miss unexpected insights could require redesigning survey techniques and feedback channels along with constant manual efforts. For organizations with unique words, misspellings, and slang, this process can feel insurmountable.



Worried that you're missing out on critical information? Look for a generalized text analytics solution. A generalized solution is one that works across all domains, use cases, and industries. These types of products have a conceptual, native understanding of the languages they support, meaning they can analyze datasets based on context and understand words and phrases they've never seen before.

Factors to consider:

- The number of languages the solution natively supports
- The time required for a user to start performing analyses
- Whether or not the solution requires libraries or ontologies
- If the solution understands misspellings, slang, and domain-specific language
- The solution's ability to understand words and phrases it's never seen before

4. What features does this product offer?

Analytics solutions aren't valuable unless they surface insights that inform business decisions. When choosing a text analytics application, consider one that allows users at all levels to elevate findings throughout their discovery journey.

The best text analytics solutions offer features and tools that help users visualize, organize, and make sense of what they've found in their data. Depending on your needs, a variety of tools may be essential to unlocking the full power of text analytics. Look for a solution that goes beyond simple discovery and a flashy UI – make sure it lets you engage with, learn from, and display your data in new ways across your entire organization.

Visualizations

Since most text analytics solutions are designed for a high volume of text data – and you've made it this far because you simply can't read all of it on your own – tools that help you visualize insights and patterns are important. Solutions that provide different views and dashboards are best to help you browse insights at a high level. Those that go beyond a basic word cloud or frequency visualization are even better: they'll let you process text data in ways specific to your business and needs, so you can quickly find and communicate insights to stakeholders.

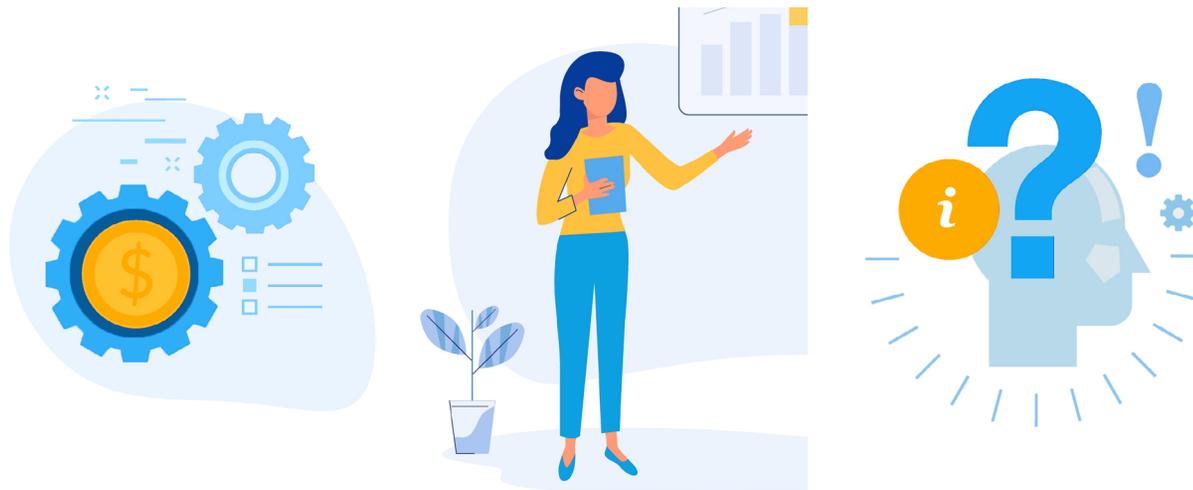
Search and filtering

A well-rounded solution needs to do more than show you high-level takeaways – you also need the ability to drill down. This helps you search keywords and phrases and cut through chatter by analyzing their importance, frequency, and relation to other terms. Then, take a closer look at relevant documents that you might not have identified without the initial analysis. A thoughtful tool sorts terms by their true relevance relative to your dataset, not just the overall language. After identifying high-level topics of interest, dig in using metadata that isolates relevant topics within subsets.



Sentiment detection

Sentiment features help organizations identify topics from a specific document and determine whether the opinion expressed in the document is positive or negative. Decision-makers can then use this data to develop a more in-depth understanding of their target audience. Sentiment helps you gauge how people feel about your offerings at a glance. Better yet – find a vendor that offers sentiment analysis at the aspect or concept level, allowing granular insight into how people really feel.



Mixed data analytics

Your data isn't only arriving in text form. You have scores and ratings attached to a lot of what you're analyzing, and you don't just want to lose all that context. Make sure your tool offers a way to analyze unstructured reviews and survey feedback, revealing how your unstructured data correlates with quantitative ratings. It's the fastest way to reveal which key concepts mentioned in your data have the greatest impact on scores. That way, business leaders can gain insights more frequently and take action more quickly.

Factors to consider:

- The ways the solution intelligently helps users process their data
- If the types of analytics features match your analysis needs
- Whether or not the solution's features require a data science background to understand

Ask, evaluate, choose. Then, get started.

Congratulations! Simply deciding to invest in text analytics means you're on your way to a faster, more effective method of understanding feedback. And by pairing your own intimate organizational knowledge with the right questions to ask vendors, you're ensuring the best outcome for your team.

At Luminoso, our generalized solutions analyze all types of unstructured text, like support tickets, open-ended survey responses, and reviews. Based on a pre-trained NLU system, Luminoso analyzes text natively in 15 languages with QuickLearn®, our proprietary natural language modeling system that automatically delivers domain-specific results based on your specific terminology. Using Luminoso's unique combination of contextual learning and background knowledge, receive intelligent insights for any dataset – no training, coding, or libraries required.

Ready to get started? [See a free demo today.](#)



The text analytics buyer checklist

1. How much time do I need for setup and maintenance?

- The amount of manual work the solution requires before machine learning can take place
- Whether or not your organization needs a dedicated data scientist to use the system
- How much specialized code the solution requires
- If business users need technical training to easily navigate the application

2. What kind of data do I need, and how much of it?

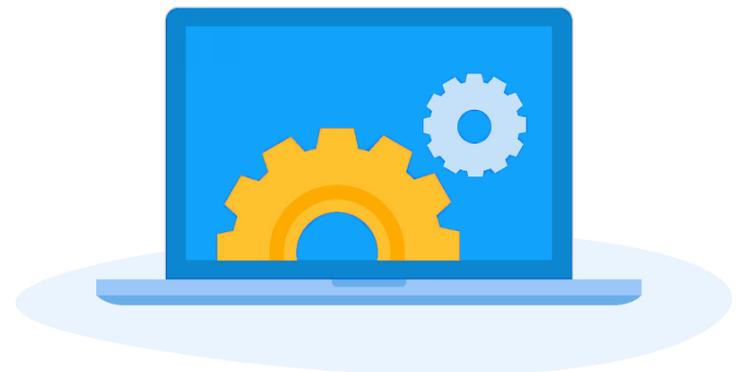
- The number of documents needed to train the system before it can understand domain-specific language
- Whether the solution's pricing is volume-based, and if so, what tiers are available
- The time and effort it takes go from a raw data upload to meaningful insights

3. How powerful and sophisticated are the insights I'll get?

- The number of languages the solution natively supports
- The time required for a user to start performing analyses
- Whether or not the solution requires libraries or ontologies
- If the solution understands misspellings, slang, and domain-specific language
- The solution's ability to understand words and phrases it's never seen before

4. What features does the product offer?

- The ways the solution intelligently helps users process their data
- If the types of analytics features match your analysis needs
- Whether or not the solution's features require a data science background



About Luminoso

Luminoso builds text analytics products for analyzing conversational text data like support tickets, open-ended survey responses, and reviews. Using common-sense artificial intelligence to understand language, we empower organizations to discover, interpret, and act on what people are telling them. Requiring little setup, maintenance, or training, Luminoso combines world-leading natural language understanding technology with a vast knowledge base to learn words from context – like humans do – and instantly analyze text without libraries. Our products provide native support in 15 languages, so leaders can explore text data, make sense of feedback, and surface business-critical insights. Luminoso is privately held and headquartered in Boston, MA.

To learn more, visit us at luminoso.com.



LUMINOSO



www.luminoso.com



info@luminoso.com



[@LuminosoInsight](https://twitter.com/LuminosoInsight)