

Solution Brief

Artificial Intelligence
Natural Language Processing



Unlocking Critical Business Insights and Enabling Data-Driven Decision Making with IBM Watson Natural Language Understanding (NLU)

IBM Watson Natural Language Understanding, powered by Intel processors and optimized with Intel software tools, uses deep learning techniques to extract meaning and metadata from unstructured data.



Cost Savings

\$6.13

million in benefits over three years¹

ROI

383% ROI

over 3 years¹

Time Savings

50%

reduction in time spent on information-gathering tasks¹

Revenue

5% annual

increase in revenue¹

Organizations are looking for new ways to extract valuable insights from their unstructured data

In today's data-driven world, enterprises are generating and receiving an ever-increasing volume of data, but unfortunately most organizations are unable to take advantage of that data and its insights. An estimated 80% to 90% of data collected is unstructured or 'dark data,' meaning it is not mapped to a predefined structure and therefore not easily searchable or accessible. And this surge of unstructured data is growing rapidly—many times faster than the rate of growth for structured databases, leaving most businesses with a significant untapped resource²

These unstructured data sources are rich with valuable information that, if made accessible, has the potential to help enterprises dramatically improve their products, services, and customer experiences. However, extracting insights from this data is comes with challenges:

- **Human error:** Manual analysis of data is tedious and repetitive, which can lead to mistakes such as misinterpreting the meaning of a sentence or missing vital information resulting in inaccurate insights and misinformed decisions.
- **Resource intensive:** Analyzing such a large volume of data is extremely time consuming, requiring significant effort and human resources to make a dent in the sheer volume of data requiring review. To fully staff this effort would require tremendous investment, most likely out of reach for most organizations to fund, not to mention the unlikelihood of humans being able to fully understand each piece of data and collate that information into meaningful insights.
- **Nuance:** Unstructured data is often written in "natural language", which can be complex and ambiguous and can change depending on context, tone, and intent. Understanding the nuances of language and identifying the key information from unstructured data can be difficult for humans, let alone at scale.
- **Required expertise:** Other solutions that attempt to automate analysis of unstructured data often require a large team of analysts or data scientists to create models that can review and interpret the data. These resources require a significant, and likely prohibitively, expensive investment for businesses.

These challenges prevent enterprises from making the most out of their data, representing a significant opportunity loss for insights that could be a catalyst for increased revenue, enhanced customer experiences, and improved operational efficiency.

To address these challenges, businesses are turning to innovative technologies to help analyze and make sense of their unstructured data. AI and natural language processing (NLP) technologies like IBM's Watson Natural Language Understanding (NLU) extract critical insights quickly and accurately from unstructured data sources, enabling enterprises to confidently make data-driven decisions.

IBM Watson Natural Language Understanding (NLU)

IBM's Watson NLU solution leverages AI and deep learning models to extract context, meaning, and metadata from unstructured text data at scale. It helps users get underneath their data using text analytics to extract business-specific categories, classification, entities, keywords, sentiment, emotion, relations, and syntax. These text analytics features surface real-time actionable insights to provide employees with the necessary tools to pull metadata and patterns from troves of unstructured data. With this new level of understanding, users can extract the meaning behind their data, including key business drivers like customer preferences and opportunities for organizational optimization and improvement.

Its powerful capabilities and domain customization make it relevant across a wide range of industries, business types, and use cases. IBM Watson NLU can be customized to meet the specific needs of a business, trained on a specific domain or industry, or customized to recognize a business' unique terminology. It can also be deployed instantly behind any firewall or on any cloud, providing flexibility on how and when the technology is used. Its valuable, business-specific insights help with use cases such as content recommendation, advertising optimization, audience segmentation, and voice of the customer analysis.

What are NLP and NLU?

Enterprises can leverage natural language processing (NLP) solutions to understand their unstructured data and make faster and smarter decisions.

1 Natural language processing (NLP)

Converts unstructured language data into a structured data format to enable machines to understand speech and text and formulate relevant, contextual responses.³

2 Natural language understanding (NLU)

A subset of natural language processing, which uses syntactic and semantic analysis of text and speech to determine the meaning of a sentence. Syntax refers to the grammatical structure of a sentence, while semantics alludes to its intended meaning. NLU also establishes a relevant data structure which specifies the relationships between words and phrases. While humans naturally do this in conversation, the combination of these analyses enables machines to understand the intended meaning of different texts.³



Four Common Watson NLU Use Cases



Content recommendation

IBM Watson NLU analyzes the text of articles, blog posts, and other content to understand overarching topics, themes, and general sentiment. These insights can help identify and recommend related content to users based on their interests and preferences.



Advertising optimization

Customer reviews, social media posts, and other user-generated content is analyzed by IBM Watson NLU to understand a content author's sentiment towards products and brands. Business decision makers can use these insights to develop advertising campaigns that target specific audiences with resonant messaging.



Audience segmentation

Customer feedback or other sources of unstructured data are analyzed for insights to help segment audiences based on their interests, preferences, and general sentiment. This information can be used to provide personalized marketing content to different audience segments based on its insights.



Voice of the customer analysis

IBM Watson NLU analyzes customer feedback, surveys, and other sources of unstructured data to understand the sentiment of feedback and identify themes that are most important to customers. This information can be used to improve products and services, identify areas for improvement, and better understand customer needs and preferences.

How the Solution Works

Input: Users start by providing a piece of unstructured text data to analyze, such as a news article, social media post, email, or other text-based document to Watson NLU via API.

Text Analysis: Then, the text is automatically analyzed by Watson NLU on IBM Cloud, powered by 4th Gen Intel® Xeon® processors, using various natural language processing techniques to help identify the meaning of individual words and the relationships between them. This can include sentiment analysis, keyword extraction, emotion detection, and more. See the table below for more specific details about the text analytics features.

Output: Watson NLU provides a structured output in standard text-based format (JSON) that includes the results of the various analyses performed on the input text that can then be used to create visualizations, automate workflows, or power chatbots and virtual assistants.

Key Benefits



Surface real-time actionable insights and give employees the tools to process and extract valuable insights from unstructured data



Improve efficiency and accuracy of data analysis, enabling data-driven decision making














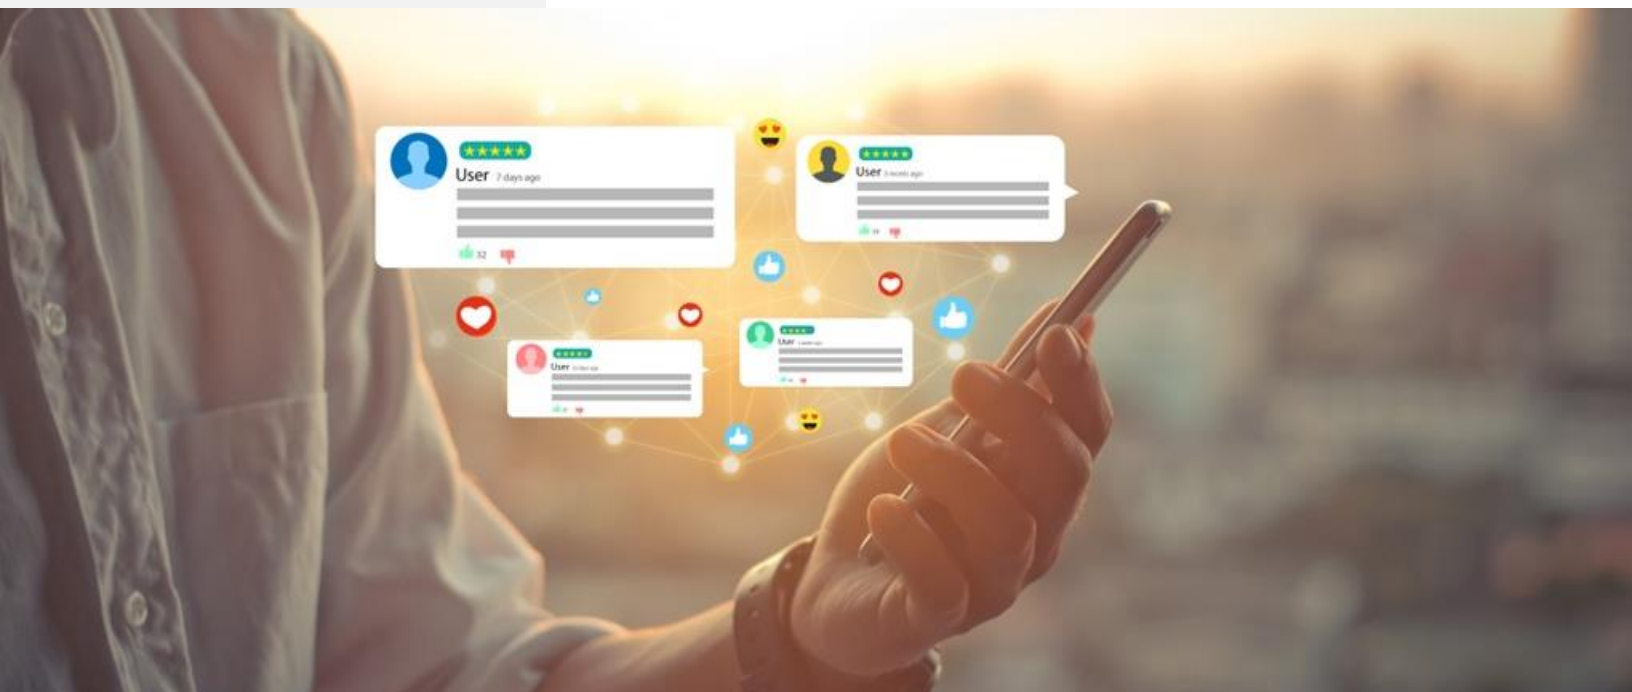
Reduce costs and risks associated with time-consuming and resource-intensive manual analysis of data



Leverage an accessible NLU solution that accommodates the needs of any business at scale—from small business to multi-national organization

Text Analytics Features

Feature	Description
 Categories	Categorize data with granularity using a five-level classification hierarchy
 Classifications	Classify text with custom labels to automate workflows, extract insights, and improve search and discovery
 Concepts	Identify high-level concepts that aren't necessarily directly referenced in the input content
 Emotion	Extract emotions (joy, anger, sadness, fear, and other feelings) conveyed by specific target phrases or by the document as a whole
 Entities	Detect people, places, events, and other types of entities mentioned in the content using out-of-the-box capabilities
 Keywords	Identify keywords from the input text and understand their relevance
 Metadata	Quickly extract information from a document such as author, title, images, and publication dates
 Relations	Understand the relationship between two entities within the content and identify the type of relation
 Semantic Roles	Parse sentences into subject-action-object form and identify entities and keywords that are subjects or objects of an action
 Sentiment (Beta)	Analyze the sentiment (positive, negative, or neutral) towards specific target phrases and of the document as a whole
 Syntax	Analyze the syntax of specific words and sentences by part-of-speech and individual character location



Intel® Technology helps Watson NLU Enhance Performance

Watson NLU leverages Intel® Technology to maximize performance and scalability for power-intensive AI workloads. Watson NLU utilizes IBM Cloud powered by 4th Generation Intel® Xeon® processors, optimizes NLU model performance with the Intel® oneAPI Analytics Toolkit, and utilizes both Intel® oneAPI Deep Neural Network Library and the Intel® Optimization for TensorFlow powered by OneAPI.

4th Gen Intel® Xeon® Processors: AI processing is a notoriously expensive and power-intensive task that requires high-performance computing power to handle the large volume of data involved. Purpose-built for AI, 4th Gen Intel Xeon Processors are equipped with up to 60 powerful cores, support a range of power features, and are optimized specifically for running complex, intensive workloads. The hardware optimizations allow Watson NLU, which runs on IBM cloud powered by 4th Gen Intel Xeon Processors, to leverage out-of-the-box software features from popular Machine Learning and Deep Learning frameworks that boost performance and throughput for demanding computational tasks like modeling and simulation, data analytics, and machine learning.

Intel® oneAPI Analytics Toolkit: This toolkit gives data scientists, researchers, and AI developers tools and frameworks to accelerate machine learning and data analytics pipelines and improve performance. It includes optimized libraries and tools that can be used in conjunction with popular machine learning frameworks. By incorporating the Intel® One API library, IBM's NLU technology takes advantage of the new hardware features and accelerators on the Intel® Xeon® processor-based infrastructure that can help improve performance.

Intel® oneAPI Deep Neural Network Library (oneDNN): Part of the Intel® oneAPI Analytics Toolkit, this open-source, high-performance library for deep learning applications is gives data scientists, researchers, and AI developers tools for building neural networks. It includes access to existing deep learning frameworks and provides optimized implementations of operations such as convolution, pooling, and normalization, which are commonly used to build neural networks. With oneDNN, Watson NLU is able to perform large-scale neural network computations efficiently, enabling the solution to analyze text data in real-time.

Intel® Optimization for TensorFlow*: Intel Optimization for TensorFlow* (available as part of the Intel® oneAPI Analytics Toolkit) is a set of software optimizations designed to improve the performance of TensorFlow*, which is a popular open-source framework for machine learning. By using Intel Optimization for TensorFlow*, Watson NLU has benefited from improved training and performance for its deep learning models on Intel hardware.

Watson NLU in action

Challenge: A technology company that provides AI solutions to legal professionals to help automate routine legal work turned to Watson NLU to address a critical pain point for their customers. They were looking to help legal organizations reduce the amount of time and resources they were investing to manually review and categorize large volumes of documents and enable them to shift limited resources to focus on higher-value tasks.

Solution: This AI technology company turned to IBM Watson NLU to build a domain-specific model focused on legal terminology and concepts to automate the tedious process of reviewing and categorizing legal documents.

They launched a highly effective, intelligent model that, combined with their own proprietary software, allows legal professionals to quickly and accurately analyze complaints and create early-phase response drafts.

Results: This AI technology company's customers are seeing impressive results after integrating Watson NLU



Legal professionals can draft early-phase documentation in **two minutes or less** – a huge improvement from current processes that take an associate a full day of work to draft⁴



Shifting the workload to the application drives down costs significantly – estimates calculate that users are **reducing costs by approximately 80%**⁴

Conclusion

IBM Watson Natural Language Understanding (NLU) offers a compelling solution for businesses seeking to leverage powerful insights from all of their text data, structured or not. With its robust natural language processing capabilities powered by Intel Xeon® processors, Watson NLU can quickly sift through data and extract valuable insights to businesses with speed and accuracy previously impossible through manual analysis. By leveraging the power of IBM Watson NLU, businesses can stay competitive in today's ever-changing, data-driven digital landscape.

Learn More

1. [IBM Website](#)
2. [IBM Watson Natural Language Understanding Solution Page](#)
3. [IBM Watson NLU Documentation](#)
4. [Improving Watson NLP Performance](#)
5. [IBM Watson NLP Performance with Intel Optimizations White Paper](#)
6. [IBM Watson NLU Community blog on Intel optimizations](#)
7. [Watson NLU Solution Demo](#)
8. [IBM Research blog on value of Intel optimizations](#)
9. [4th Gen Intel® Xeon® Processors](#)
10. [Intel® oneAPI Analytics Toolkit](#)
11. [Intel® oneAPI Deep Neural Network Library \(oneDNN\)](#)
12. [Intel® Optimization for TensorFlow*](#)



About IBM

The International Business Machines Corporation (IBM), is a multinational technology corporation headquartered in Armonk, New York. Founded in 1911, IBM is one of the oldest technology companies in the world and has a long history of innovation and leadership in the industry. IBM works to integrate technology and expertise, provide infrastructure, software (including market-leading Red Hat) and consulting services for clients as they pursue the digital transformation of the world's mission-critical businesses.



Sources

1. [The Total Economic Impact™ Of IBM Watson Natural Language Processing \(NLP\) Solutions](#), Forrester, 2021
2. [Tapping the Power of Unstructured Data](#), MIT Management Sloan School, 2021
3. [NLP vs. NLU vs. NLG: the differences between three natural language processing concepts](#), IBM, 2020
4. [IBM Website Case Study](#), IBM, 2021

Notices & Disclaimers

Intel is committed to respecting human rights and avoiding complicity in human rights abuses.

See Intel's [Global Human Rights Principles](#). Intel® products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

Intel technologies may require enabled hardware, software or service activation. No product or component can be absolutely secure. Your costs and results may vary. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.