The 2023 Expert NLP Survey Report

Trends driving NLP Investment and Innovation

expert.ai

Introduction

A previous expert.ai report, Harnessing the Power of Unstructured Data with NLP and NLU, revealed that organizations are keenly aware of the value of data to their operations, but those not utilizing NLP and other text analytics tools may be failing to exploit critical data assets.

A recent Forrester Research report highlighted that companies driving decisions using advanced data analytics are eight times more likely than data analytics newbies to have a revenue growth rate of 20% or more. Customer communications, emails, contracts, reports, and digital archives are all potentially valuable assets for language-based process automation and improvement.

Natural language processing (NLP) is the key that unlocks the potential of AI to comprehend and utilize unstructured language data, bridging the automation gap between humans and technology and leveraging existing assets for new insights that were previously unavailable. Its uptake is being driven by big data, digital transformation, and the rise in human-machine communications. Its use cases stretch across all business operations, from marketing to finance, customer care to sales.

For companies that have been employing NLP solutions for several years, the most compelling arguments for the technology are improvements in efficiency, reducing risk and, above all, cost reduction. AI has become embedded in their business models, and NLP is allowing them to create new opportunities.

NLP budgets are increasing, and the value of the market is exploding, with Fortune Business Insights predicting that it will be worth $127 billion by 2028. Our NLP practitioner research shows that four in five organizations surveyed have NLP models in production, with tens of thousands of documents being processed by each monthly.

But organizations still face challenges aligning with stakeholders over which use cases to prioritize, and in justifying the costs associated with NLP modeling and tools. Without a clear understanding of the varying attributes of alternative technologies, many are failing to exploit the full capability of NLP. Because of these factors, they risk fumbling the competitive advantage that NLP can deliver.

At expert.ai, we have 30 years of experience in developing natural language solutions. Our natural language understanding (NLU) platform offers NLP value to all types of organizations, from those just beginning to investigate the possibilities of NLP to experienced users with multiple solutions in production across the business.

This report reveals where each of these groups stands, the solutions they're employing, challenges they face, their expectations, how they track return on investment, and future NLP spending predictions. By mapping current NLP adoption,we gain clarity of what organizations require from NLP solutions, both now and in the future.

Marco Varone

Founder & Chief Technology Officer, ***expert.ai***

State of NLP in 2023*

Increased Spending

77%

of organizations surveyed expect to spend more on NLP projects in the next 12-18 months

Platform Preference

71%

use a natural language platform that supports multiple text analytics use cases

Accuracy Challenges

67%

of businesses with NLP models in production for 5+ years still deal with accuracy challenges

ROI Efficiencies

61%

of NA teams measure ROI based on improvements in efficiency / efficiency improvements

Responsible AI

74%

74% say they consider how responsible (explainable, energy efficient and unbiased) each AI approach is when selecting their solution

Investment in NLP is on the rise… our results show companies have a strong preference for: platforms that support multiple use cases, especially those with strong model development and maintenance capabilities; use cases that improve efficiency and help them get to market faster; and allow them to meet their core challenge of delivering highly accurate results.

*Based on research among current NLP practitioners across business and technical roles in North America and Europe.

Executive summary

With digital transformation driving a new imperative for processing unstructured data across multiple industries, artificial intelligence solutions incorporating natural language processing (NLP) are becoming business critical.

While some sectors are more advanced in the implementation of NLP, this report reveals an uptick in interest in natural language projects, with more being incorporated across various business operations. Prepared by the AI Journal from research among current NLP practitioners across business and technical roles in North America and Europe, the report highlights the key benefits businesses are seeking and the types of solutions currently being employed.

Acknowledging the ways in which NLP has yet to evolve, this report also addresses the challenges faced by companies looking to initiate such models, how to track the return on investment, and the anticipated growth in NLP spend and implementation in the coming years.

Emerging trends surfaced by the research include:

Where and how NLP is being deployed
Four out of five (80%) of organizations surveyed have NLP models in production.

Most NLP projects are relatively new, with 47% in production for less than two years. Those businesses still experimenting with NLP models make up 19% of respondents.

Flagged by 47% of respondents, data protection and governance (i.e. GDPR and Personal Identifiable Information) use cases are the main business areas that teams are currently using or considering using NLP solutions to overcome.

Extraction and named entity recognition (NER) projects are the top technical use case, at 47%, where businesses are looking to automate their language-intensive processes.

Core challenges in NLP deployment
The top two challenges faced by businesses when adopting an NLP solution are aligning with various stakeholders on which cases to prioritize (38%), and the costs associated with NLP modelling and tools (38%).

Where an organization is on its NLP development journey affects the main challenges it faces. Of those still evaluating NLP use cases, it is data security and governance (64%). Businesses starting to experiment and build NLP models cite choosing which AI approach to use to achieve the desired model results (47%). Those with NLP models in production for less than two years flag aligning with stakeholders on which NLP use cases to prioritize (45%). For those which are two to four years into production, the main issue is building the business case for NLP projects (42%). And two-thirds (67%) of businesses with five-plus years of NLP models said their biggest challenge was achieving the level of accuracy or quality needed to put those models into production.

Methodology matters: Combined solutions versus standalone
For NLP use cases currently in production, 52% of respondents use a mixed machine learning and symbolic/rules-based approach, while 79% employ machine learning with either deep learning or graph AI.

Organizations are using combined solutions to create more robust outcomes, exploiting the attributes of each technology to produce higher quality, more accurate results. The majority of organizations are using some form of hybrid (aka composite) approach, with just 6% employing machine learning on its own.

Platform versus cloud
71% of respondents describe their technology solution for NLP use cases as a natural language platform that supports multiple text analytics use cases.

The most common reason for choosing natural language platforms are their model development and maintenance capabilities (47%).

Those selecting cloud solutions often do so because of the support for machine learning AI (45%), while many of those using open-source NLP software appreciate the natural language understanding capabilities (44%). A significant proportion of users (40%) value case-specific SaaS point solutions for their technical performance and scalability.

Tracking ROI is a work in progress
More than half of organizations measure the return on investment for NLP projects based on one or more of the following metrics: time to production (54%), efficiency improvements (53%) and cost reduction (53%).

Efficiency improvements carry greater weight depending on geography, with 61% of businesses in North America employing it as a measure, compared to 36% in Europe.

Other common ROI metrics include competitive advantage, reduction of risk and increase in revenue.

Strong growth predictions for NLP spending
77% of respondents expect their spending on NLP projects to increase over the next 12 to 18 months.

Very few of those surveyed predict a decrease in spending on NLP capabilities, and only 13% expect spending to stay the same. Of those anticipating an increase, 39% think they will see spending rise by up to 10%, and 35% anticipate it increasing by between 11% and 20%.

Correlation can be found between spending predictions and how mature an organization is in its NLP journey.

NLP usage:
Current and predicted deployment

According to IBM's Global AI Adoption Index 2022, 35% of companies worldwide are using artificial intelligence (AI), with an additional 42% exploring the potential to incorporate it into their operations.

As part of the growing implementation of AI, NLP is gaining greater traction, and is being deployed in a diverse and expanding number of settings. Most frequently seen in customer care applications (38%), IBM reports that security operations (36%), business development (32%), sales, marketing, human resources and finance are all fields where NLP is helping to automate language-heavy processes.

While symbolic/rules-based approaches to NLP have been around for 50-plus years, the application of machine learning approaches to NLP is still in its relative infancy – only 8% of respondents from our research have had NLP models in production for more than five years, while 25% have been working with them for just two to four years. Fortune Business Insights predicts that the NLP market will grow from $21 billion in 2021 to reach $127 billion in 2028.

It comes as no surprise, therefore, that four in five (80%) of the organizations we surveyed have NLP models in production, although some industries are further ahead in the adoption journey. We see that the vast majority (95%) of respondents in the financial, banking and insurance sector are employing NLP models, compared to 73% in the technology and software sector.

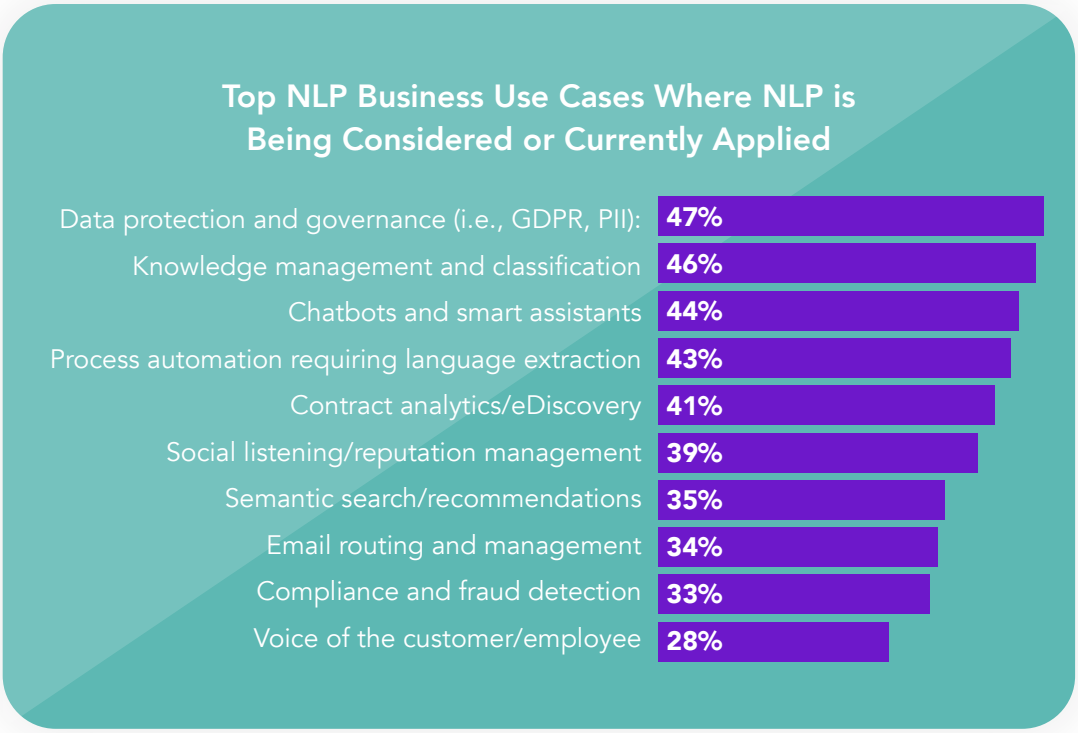

Among organizations currently using or considering using NLP solutions, 47% are looking to it to combat data protection and governance challenges. This figure rises to 50% in North America, while it is a concern for 38% of European businesses.

Other functional business issues where NLP is providing solutions include knowledge management and classification (46%), chatbots and smart assistants (44%), and process automation (i.e., claims automation in insurance) that requires language extraction (43%).

The main NLP technical use case applied to these documents is extraction or named entity recognition, with 47% of businesses considering or currently using it (50% in North America and 38% in Europe). Other popular use cases include intent classification/ chatbots (44%) and natural language generation (44%).

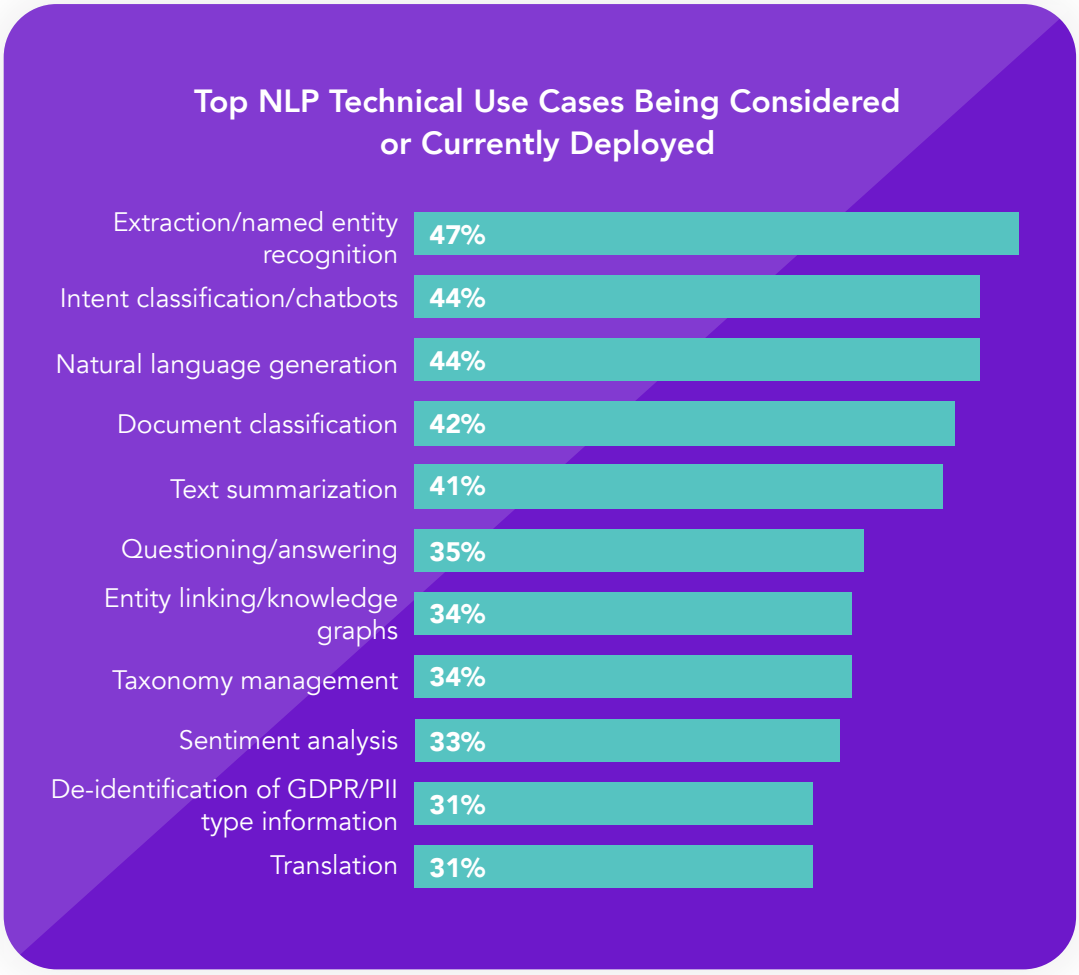

NLP challenges: Issues faced by businesses adopting natural language models

As with the adoption of any rapidly evolving technology, organizations face several significant barriers to the development and implementation of NLP models. Our data shows that only 1% of current NLP practitioners report encountering no challenges in its adoption, with many having to tackle unexpected hurdles along the way.

> *"Being overwhelmed at the beginning made our company take a lot of wrong turns."*

In its 2022 Global AI Adoption Index, IBM reports that natural language capabilities are generally perceived by organizations as expensive to implement. Over half (54%) of IT professionals cited costs as a large or medium barrier to introducing NLP.

In line with this, our research found the financial cost associated with NLP modelling and tools was the largest challenge faced when adopting the technology, identified by 38% of respondents.

> *"Stakeholder buy-in is the main challenge."*

Today, organizations are concerned with the priorities of all of their stakeholders and finding alignment between them can be demanding. It's little surprise, therefore, that in the realm of AI, getting all business stakeholders to agree on the appropriate usage of NLP models is proving difficult for 38% of our respondents. As familiarity with NLP attributes increases and stakeholders get on board with the benefits, these alignment obstacles should start to clear.

The greatest influencer of stakeholder support for NLP adoption is likely to be the perceived accuracy of results, currently an issue for 37% of organizations, according to our figures.

> *"How AI deals with colloquialisms and slang words used by employees is an issue."*

While continued monitoring and modification will increase the accuracy of NLP models in the future, the evolution of language terms and data drift will always tend to change over time. It is unlikely that users will ever feel confident enough to stop improving the quality of the output, or to remove expert human oversight.

> *"It's hard to find people who can do these jobs around language processing."*

Locating NLP expertise isn't always easy, and 29% of our respondents cite finding employees with the skills to use NLP tools effectively as a hurdle to implementation. With the war for talent raging, this issue is unlikely to resolve in the near future.

"Building the business case and ROI for NLP."

Building the business case for NLP projects, especially in terms of return on investment, is another major challenge facing would-be users – raised by 37% of businesses. Later in this report, we will address the ways that organizations are making their case for NLP models, in terms of time to production, cost reductions and improvements in efficiencies.

NLP maturity matters

Our data reveals that the maturity of an organization's NLP approach dictates the nature of the challenges it faces.

Top 3 Challenges by NLP Maturity Level

Evaluation and experimentation phase (no NLP models in production)	Early phase NLP implementation (under two years in production)	More mature NLP models (two years-plus in production)
Data security and governance	Aligning with business stakeholders on which NLP uses cases to prioritizegovernance	Building the business case/ ROI for NLP projects
Building the business case/ ROI for NLP projects	Data security and governance	Achieving the level of accuracy/quality needed to put NLP models in production
Choosing which artificial intelligence approaches to use to achieve the desired model result	Costs associated with NLP modeling and tools	Costs associated with NLP modeling and tools

Of those still evaluating NLP use cases, the biggest challenge is data security and governance (64%). Their main focus is on security and compliance and justifying the business value of having a use case, before figuring out how to build the AI. Meanwhile, businesses starting to experiment with and build NLP models struggle most with which AI approach to use to achieve the desired model accuracy and results (47%).

With a little more experience, teams are more focused on collaborating with the business on where to go next. Teams with NLP models in production for less than two years cite aligning with stakeholders on which NLP use cases to prioritize (45%). Next, they are looking at how to work more cost effectively.

More experienced teams are building business cases to justify their activities and driving high quality results to support ROI. For those with two to four years with models in production, the main issue is building the business case for NLP projects (42%). Two-thirds (67%) of businesses with five or more years of experience employing NLP models said their biggest challenge was achieving the level of accuracy or quality needed to put those models into production. These teams then move to focus on managing costs and building repeatability.

NLP technology: The most popular methodologies

When implementing NLP models, methodology matters. We asked respondents about the type of technology solutions they were using, whether machine learning (ML), symbolic/rules-based approaches, deep learning (DL) or graph – or a combination of these solutions.

For NLP use cases currently in production within organizations, we see many businesses taking a combined approach. Half (52%) use a mixed ML and symbolic/rules-based approach, while 79% employ ML with either DL or graph technology. Compared to this, only 6% use ML alone.

Definitions of NLP methodologies

Machine learning (ML): The use of computer algorithms that build a model based on sample or 'training' data, to make predictions or decisions without being explicitly programmed to do so.

Deep learning (DL): A class of machine learning algorithms that can be applied to the world of structured or unstructured data. DL is part of a broader family of machine learning methods based on artificial neural networks with representation learning to classify concepts from images, text or audio.

Symbolic/rules-based: A symbolic methodology is an approach to developing AI systems for NLP based on a deterministic, conditional approach. A symbolic approach designs a system using very specific, narrow instructions that guarantee the recognition of a linguistic pattern.

Graph: A knowledge graph is a graph of concepts where every concept is linked to at least one other concept, and the quality of this connection can belong to different classes. The interpretation of every concept is represented by its links, and every node is the concept it represents based on its position in the graph.

Hybrid: Hybrid AI is any artificial intelligence technology that combines multiple AI methodologies. In NLP, this often means that a workflow will leverage both symbolic and ML techniques.

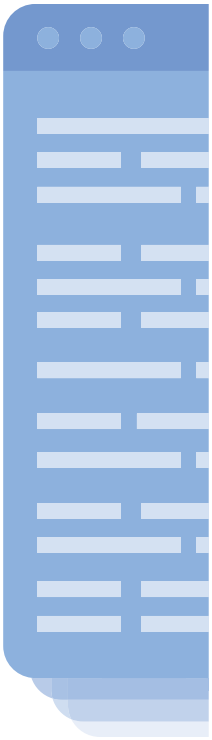

Benefits of a combined approach

Many NLP experts responsible for adopting and implementing projects in the enterprise share the strongly held belief that deep learning and machine learning are the only techniques available to develop these solutions. In doing so, they limit the areas to which NLP can be applied, as these techniques can have serious intrinsic flaws when applied alone to 'real world' NLP use cases across the business.

A major problem when implementing deep learning and machine learning techniques for NLP is that they require a significant amount of textual data for training purposes. The data available in the enterprise for complex domain specific use cases, however, is usually insufficient. There simply aren't enough relevant documents to effectively train a system, or the related privacy and data sharing issues make them unavailable or difficult to use.

In the business world, unlike the consumer world, such complexity and specificity are the very nature of enterprise documentation. This is especially true when working with complex use cases, like claims handling in insurance or contract analysis that require the support of subject matter experts (SMEs). Tasking SMEs with manual and repetitive training activities is very expensive. The menial work distracts from the high-value tasks they were hired for and often has an negative impact on employee satisfaction and engagement.

Lastly, lack of explainability is a characteristic typical of machine learning and deep learning: they are black boxes. This means that once the model has been trained, it is not possible to understand why it behaves in a certain way. This is problematic from an organizational point of view because there is no way to identify and correct any bias, or any unwanted or simply unethical behavior.

> *"Our research confirms that the best results are achieved by a combination of AI that is both knowledge-based and ML-based, hybrid AI." Forrester Research*

Machine learning and symbolic are often pitted against each other as mutually exclusive options to natural language understanding. This has forced organizations to compromise in one way or another. Recent technology advances have resulted in the availability of a hybrid approach, where organizations can use both ML and symbolic in tandem, enabling them to realize the core benefits of each. Teams often need the flexibility of a hybrid approach that integrates symbolic and ML AI techniques to achieve the success metrics most valuable to each use case, such as explainability, scalability and accuracy.

Typically using a hybrid AI approach that combines ML and symbolic can:

- Deliver highly accurate results in situations where little training data is available.
- Be much more cost effective and energy efficient than computation-heavy ML alone or neural networks.
- Improve explainability and be more predictable than ML-only approaches.
- Codify domain and process specific knowledge by including SME feedback directly in the rules and make it is easier to fine-tune results.

Which approach to take?

In response to our survey, 71% of businesses described their technology solution for NLP use cases as a natural language platform that supports multiple text analytics use cases. However, this is not set in stone, with 77% of respondents agreeing that they are encouraged to try different AI approaches when working on NLP projects. When planning these projects, 74% say they consider how responsible – that is, how explainable, energy efficient and unbiased – each AI approach is when selecting their solution.

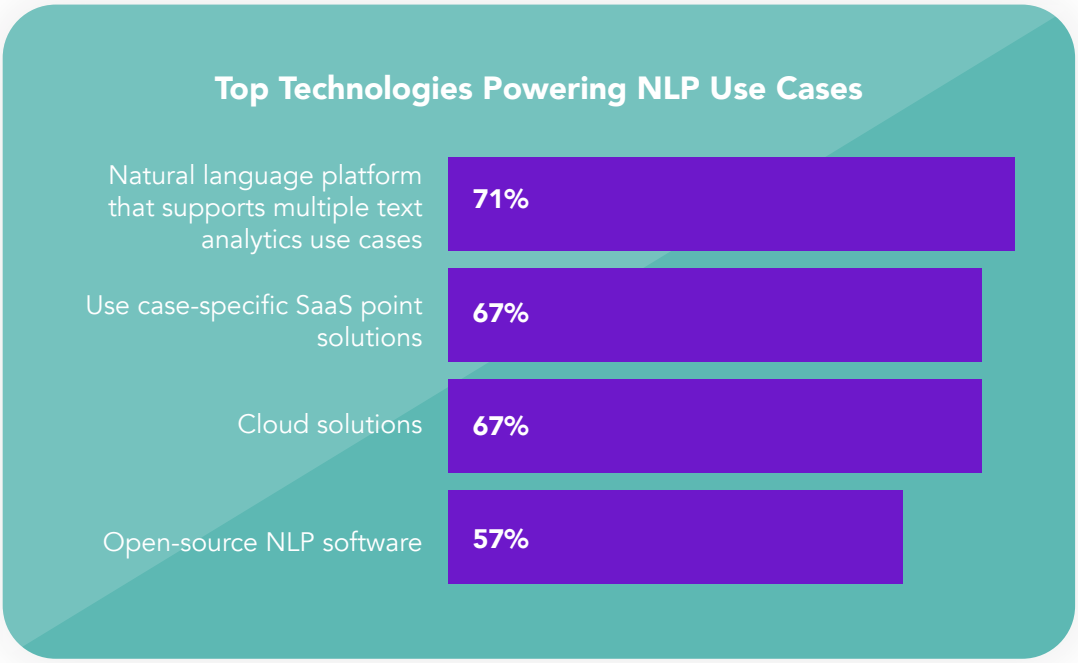

Among those using cloud solutions, our data shows that Microsoft Language API is the most commonly employed NLP services cloud provider at 75%, followed by Amazon Comprehend (70%), Google Natural Language (67%) and IBM Watson (56%).

NLP rationale: Reasons for selecting NLP software

Just as the use cases for NLP models vary – from customer care to business development and finance functions – so do the main reasons for organizations choosing different types of software, whether it's a platform, cloud or open-source solution.

Among those businesses that chose open-source software, the rationale varies depending on the attributes of the specific product, but overall, 44% selected it for process automation integration with RPA, and 44% for natural language understanding capabilities.

For cloud solutions, 45% picked them for the machine learning AI approach. Natural language platforms appealed to 47% for their model development and maintenance capabilities, while use case-specific SaaS point solutions were popular with 40% of users for their symbolic AI approach, and 40% for their technical performance and scalability.

By interrogating the top five preferred features for each solution type – cloud, open-source, platform, point solution – we can formulate an 'NLP buyer's guide' that highlights the most desirable capabilities of each.

Most Desirable Attributes by Type of NLP Solution

Cloud	Platform	Point Solution	Open Source
Machine Learning AI Approach	Model development and maintenance capabilities	Symbolic AI Approach	Natural Language Understanding capability
Out-of-the-box Knowledge Models	Process Automation Integration with RPA	Technical Performance and Scalability	Process Automation Integration with RPA
Technical Performance and Scalability	Integrated natural language workflow depth and flexibility	Multi-language support	Out-of-the-box packaged solutions
Integrated natural language workflow depth and flexibility	Ability to support multiple use cases	Natural Language Understanding support	Integrated natural language workflow depth and flexibility
Ability to support multiple use cases	Out-of-the-box packaged solutions	Model development and maintenance capabilities	Technical Performance and Scalability

From this data, we can see that open-source software is selected for its NLU capabilities, process automation potential and out of the box packaged solutions. Cloud solutions appeal for similar reasons, but they are particularly favored for their out of the box knowledge models, ability to support multiple use cases and for their machine learning approach.

Point solutions are obviously more targeted to specific use cases, and therefore preferred for their rules-based symbolic approach, multi-language support and simplified model development approach. Meanwhile, natural language platforms, utilized by 71% of our respondents, in many ways offer the best of all worlds for those not taking a combined multi-vendor approach. Support for multiple use cases, out-of-the-box solutions and model development and maintenance capabilities are all desirable features for those looking to use NLP models across multiple business units.

Builder beware!

In a recent Forrester Research report on Text Analytics Strategy, Forrester recommends, in most cases, buying a commercial solution. Even if companies could successfully build the NLP using open source, builders would still need to find the resources and budget to:

- Build a UI for business users.
- Build a UI for LangOps model orchestration and maintenance.
- Build and maintain all the data integrations.
- Monitor for performance and accuracy issues.
- Train and plan for employee succession; staffing changes

NLP ROI: How return on investment is measured

According to a 2019 survey by McKinsey, 63% of executives whose companies had implemented AI projects reported revenue increases in the relevant business units as a result, and 44% saw cost savings. But despite these glowing reviews, of the businesses surveyed for this report, 37% cited building the ROI case for NLP projects as an impediment to adopting the technology. In Europe, that figure rises to 44%.

The traditional calculation of ROI percentage – gains minus cost of investment divided by cost of investment – is made challenging by the need to estimate potential variable operating costs on as-yet undeveloped NLP solutions.

One approach to overcome this barrier is using a variety of methods to present the case for natural language processing to stakeholders, employing differing ROI metrics to track the success of existing models, and so predict the likely returns from new projects.

Our data shows that more than half of organizations measure the ROI for NLP based on one or more of the following factors: time to production (54%), efficiency improvements (53%) and cost reduction (53%).

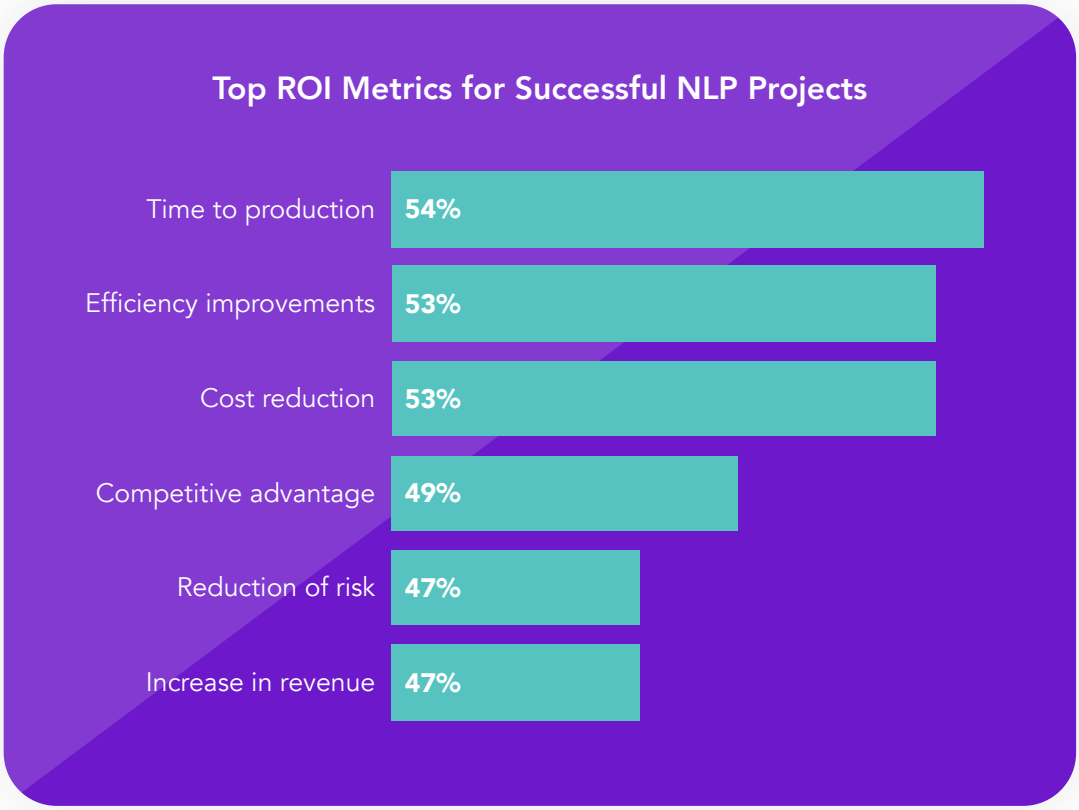

NLP budgeting: 77% expect a spending increase

Despite concerns around costs, it appears that organizations are willing to step up budgets to implement NLP. We note from our data that more than three-quarters (77%) of businesses are optimistic about NLP spend in the coming 12 to 18 months. Meanwhile, only 9% expect spending to decrease. This bullish approach is emphasized by the fact that, of all those expecting higher spend on NLP projects, almost half anticipate that increase to be by 11% or more.

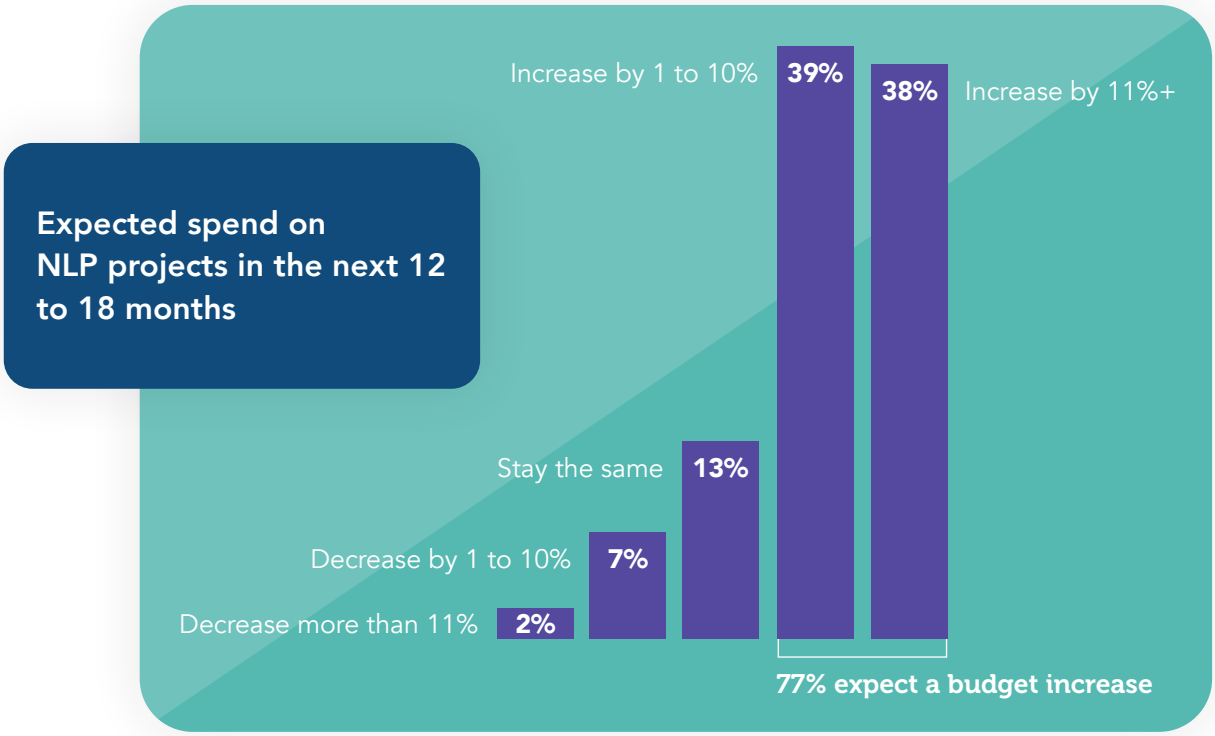

Predictions of greater spend are strongest among NLP users with more mature models – defined as in production for more than two years – with 86% anticipating greater investment in natural language processing.

Two-fifths of these organizations say their spend increase will be more than 10%, reflecting how embedded NLP has become in the business operations of these early adopters.

Among those in the early phases of NLP implementation – organizations that have had models in production for less than two years – there is also optimism over spend. This group makes up almost half (47%) of those expecting to increase spend by up to 10%, and 46% of those expecting it to rise by 11% or more. Overall, more than three-quarters (76%) of recent adopters anticipate greater investment in NLP.

Organizations that are still evaluating NLP use cases and experimenting with and building NLP models are more cautious in their forecasting, but their investment expectations are still appreciably optimistic. Two-thirds (64%) are planning to increase spend, with 39% predicting that increase will be in excess of 10%.

Conclusion: What does this tell us?

This report identifies current themes outlining what businesses are hoping to achieve with NLP solutions, how return on investment is quantified, and the types of solutions being employed. It's clear there is a substantial appetite for operational efficiencies to reduce costs, drive growth and gain a competitive advantage delivered by using NLP, with only 1% of our respondents not actively considering NLP business use cases, and 77% expecting to increase NLP spending.

Hybrid approaches work best
The vast majority of companies surveyed for this report are taking a multi-solution approach to NLP. Most NLP projects will benefit from the flexibility of a hybrid AI platform that integrates symbolic and machine learning AI techniques to achieve the success metrics most valuable to each use case, such as explainability, scalability and accuracy.

Challenges will be overcome
As NLP becomes more ubiquitous, staffing skill levels will increase and existing models will be shown to increase efficiency, reduce risk, deliver competitive advantage and cut costs. In addition, the challenges of aligning stakeholders, demonstrating ROI and perceived cost barriers will recede, and organizations will move to rapidly implement NLP solutions.

Businesses know what they want to achieve
Four in five (81%) of respondents said their company has a clear vision of how to use AI natural language projects, with 34% strongly agreeing to the statement. Despite the potentially endless range of potential use cases, organizations already have clarity on what they want NLP to deliver. Companies should hedge against upcoming AI regulatory requirements for transparency/ explainability and ensure that they adopt responsible AI approaches sooner rather than later.

There's ambition for the future
Across the board, no matter what stage they are at on their NLP journey, businesses are overwhelmingly looking to increase investment – and grow the number of NLP models in production over the next 12 to 18 months. With such rapid expansion will come the need for deeper understanding of which technology can deliver the most accurate, repeatable and responsible approaches.

Survey Methodology

The survey was conducted among 150 NLP practitioners with active or planned projects. Research took place across the USA and Europe and the interviews were conducted online by Sapio Research in 2022."

Get Started

Interested in learning how NLU AI will transform your company? Get started here.

See what expert.ai can do for you!

expert.ai

About us

Expert.ai (EXAI:IM) is a leading company in AI-based natural language software. Organizations in insurance, banking and finance, publishing, media and defense all rely on expert.ai to turn language into data, analyze and understand complex documents, accelerate intelligent process automation and improve decision making. Expert.ai's purpose-built natural language platform pairs simple and powerful tools with a proven hybrid AI approach that combines symbolic and machine learning to solve real-world problems and enhance business operations at speed and scale. With offices in Europe and North America, expert.ai serves global businesses such as AXA XL, Zurich Insurance Group, Generali, The Associated Press, Bloomberg INDG, BNP Paribas, Rabobank, Gannett and EBSCO.

www.expert.ai **info@expert.ai**