

White Paper

Delivering ROI with AI: What Is an AI Insight Engine, and How Do I Prepare to Make Use of It?

Sponsored by: IBM

David Schubmehl

Donna Nitchie

April 2018

IDC OPINION

Many types of challenging and difficult questions directly impact the operations, performance, and profitability of organizations. These questions need to be answered with data that crosses departments and domains and usually require business analysts to collect and conduct research across a range of repositories and data types. There's little disagreement that something must be done differently to keep up with the insights hidden in growing volumes of data and that better use of data is increasingly differentiating companies with customers and with employees. Analysts and other knowledge workers need to connect and synthesize data from many different sources. The best are able to "connect the dots" with seemingly unrelated data. Artificial intelligence (AI) insight engines perform this type of functionality — uncovering connections — by finding, organizing, and relating information from many different sources to create actionable analysis and decision making.

The ability to generate insights using AI is no longer just nice to have; it is necessary. Yet unlocking the value in data is not a trivial task. Many organizations face roadblocks because of a lack of appropriate technology and processes or, more commonly, siloed and legacy data. While these challenges can derail less successful organizations, leaders are navigating through barriers and capitalizing on their data with AI-driven insights and enabling a new generation of knowledge worker roles to deliver new or increased value with improved efficiency. Benefits are being realized across industries such as insurance, manufacturing, finance, advertising, and transportation and from business areas such as:

- Sales and client engagement
- Support roles in technical field service and customer service
- Production innovation and research
- Marketing, including brand and campaign optimization

Organizations need to prepare today for this next generation of tools, and that includes potential changes to business process, culture, and adoption of new AI-based technologies. The impact of such a move toward a data-centric enterprise enabled by AI will drive new areas of revenue growth from enhanced product offerings, operations, and business relationships.

IN THIS WHITE PAPER

Extracting widespread meaningful insights from data is the revolutionary extension of search engines coupled with AI technologies that analyze, organize, access, and provide advisory services based on a range of structured and unstructured information. Organizations that unlock the value of this content have demonstrated an ability to increase revenue, improve productivity, reduce costs, respond to customer or stakeholder needs more quickly and accurately, and help bring products to market faster. This white paper presents considerations for organizations to capitalize on a technology empowered by a world of increasing data volume, variety, velocity, and veracity.

SITUATION OVERVIEW

The availability of a wide variety of data and the ability to leverage that data provide organizations with an opportunity to make better decisions, personalize customer interactions, optimize operations, and innovate. A big part of realizing this potential is dependent on efficient and effective access to unstructured content and the analysis of such content in addition to, and in conjunction with, structured data. The unstructured content, especially, is locked in a variety of formats, locations, and applications made up of separate repositories that don't talk to each other and that traditional technologies can't access or interpret.

AI technologies use a wide range of information access processes combined with deep learning and machine learning to provide expert assistance in the workplace. These systems use information curation, information retrieval, knowledge graphs, relevancy training, anomaly detection, and numerous other components to help workers answer questions, predict future events, provide recommendations, and take actions to fix issues. AI systems are beginning to be combined with collaborative technologies to provide the next generation of tools to get work done.

IDC expects that these collaboration tools will become ubiquitous over the next five years and that by 2021, 90% of new intelligent systems will have an embedded decision-centric computing architecture that automatically detects and evaluates conditions and makes recommendations about how to respond. This will fuel a massive reengineering of the workplace, making it more responsive, agile, and able to facilitate data-driven decision making in all areas of businesses.

Who Is Succeeding with AI?

IDC has identified the characteristics of organizations that are able to extract more value out of the information available to them. Leader organizations make use of information access and analysis technologies to facilitate information unification, location, discovery, and sharing among their employees and other stakeholders. These insight leaders are characterized by:

- Strategic use of unstructured information as a data asset
- Efficient access to unified information
- Effective search capabilities (including dashboards)
- Effective sharing and reuse of information among employees and other stakeholders
- Access to subject matter experts (SMEs)
- Effective leverage of selected unstructured data

As with many emerging technologies, critical success factors for implementation are related to people, processes, and technologies. Failures can stem from perceptions about the maturity of the technology, employee resistance, and a lack of understanding of potential benefits within companies.

Business leaders need to address all these concerns, not just the technology issues, to increase the successful application of AI implementations. Success or failure depends on how well a business incorporates cognitive/AI capabilities into its overall strategy.

To succeed from a technology perspective, AI implementations should be easy to implement, provide the ability to continuously learn and adapt to changing context, and provide measurable results. From a business process perspective, for example, AI adoption can be inhibited by governance and regulatory implications, including data sovereignty, tightened controls, and governance policies and procedures. Success will hinge on a well-considered data strategy for AI applications. IDC's *Executive Guide to Assessing Tangible and Intangible Impacts of Cognitive Computing and Artificial Intelligence* (IDC #US42348117, March 2017) offers recommendations as businesses adopt these transformational technologies.

To date, recommendation- and/or prediction-based AI applications are popular in a variety of industries and use cases. For example, in banking, these applications utilize embedded cognitive/AI computing capabilities to help businesses capture and extract detailed information and insight from enterprise documents such as processing of loan/mortgage application needs and requirements; automatic classification and understanding of the document – including format structure, words, and numeric information; and integration with recommendation systems (intelligent filtering engines that narrow the decision-making process) for the purpose of identifying best matching financial products or services.

Conversational AI applications have also proven to be popular. For example, in the retail industry, these applications provide customer service via a learning program with a conversational interface that understands customer needs and problems and helps reduce the time and resources spent in achieving customer issue resolution.

However, both are distinct from an AI-based insight engine that collects and organizes information using advanced text analytics, connects the dots between documents and structured data using tools such as knowledge graphs, and then assists the user with finding the answers to challenging questions using the relationships between subjects that might not have direct correlation but have indirect correlations that are identified by the AI-enabled solution.

Organizations are preparing for these types of AI-enabled solutions in a variety of ways, with participation that spans functional areas. Information audits provide an understanding of the types of information currently being collected and aggregated across organizations. Cross-functional teams collaborate to identify missing categories of information and make recommendations for acquiring the information through third parties. IT organizations are creating and maintaining data lakes where this information can be accessed and managed.

IDC found that organizations with the ability to unlock the hidden value of information are five times more likely than others to experience benefits that exceed expectations. In other words, leaders can unlock the value from their organization's information assets significantly more frequently. This has been done in a variety of industries and business functions. For example:

- Insurance companies are using advanced search and natural language processing to address high volumes of email and call center traffic. AI systems automate email responses, find and fill in missing information, and provide immediate and relevant information to call center agents handling customer inquiries.

- Financial and insurance organizations are using AI not only for handling claims or for fraud and compliance investigation but also for scaling sales expertise and personalizing customer experiences.
- Healthcare organizations are building diagnosis systems based on analyzing combinations of structured and unstructured data.

The following types of organizations are also starting to build prediction and recommendation systems:

- Manufacturing and industrial products companies are using AI for predictive maintenance applications, analyzing volumes and varieties of data from sensors (IoT) embedded in equipment to provide service before failures occur.
- In retail, analytics engines provide customized product recommendation systems both online and to sales advisors.

What Types of Questions Are Best Answered with AI?

The types of questions that are best answered with AI are those that require input from many different data sources and often aren't simple yes/no answers. In many cases, these types of questions rely on the concept of knowledge-centric reasoning, which involves AI making connections across data sets and using reasoning strategies to combine insights about entities and relations. This is often done by building a broad-based information repository that contains unstructured and semistructured data across a range of topics (commonly called a knowledge base) and then using a knowledge graph that supports the AI-based reasoning. A knowledge graph is a technology usually built on top of a graph database that acquires and integrates information into an ontology and applies reasoning-based algorithms and processes to derive new knowledge.

This type of approach is especially useful for problems that require sifting through potentially very large volumes of data looking for the right needle in a tractor trailer full of needles. Insight engines based on these technologies are also able to surface patterns and trends from large data sets using AI and machine learning that would otherwise remain hidden in the data.

How do these systems do this? AI systems facilitate the discovery, use, and collaboration of information in analysis and decision making. These technologies use information curation, information retrieval, knowledge graphs, relevancy training, anomaly detection, and numerous other components to help workers answer questions, predict future events, surface unseen trends, provide recommendations, and take actions to fix issues. Examples of these types of applications include:

- Expert analysis and recommendation systems
- Predictions and preventive maintenance
- Intelligent advisors that learn as circumstances, and data, change

Preparing for AI

Preparing for successful AI implementations involves two categories of activity: technology/processes and culture/socialization. The previous sections address the former, that is, the capabilities of AI technologies, processes for preparing data, and identifying the questions for the AI system. Successful organizations understand the relevance and applicability of different types of AI technology to solve insight generation problems, especially for knowledge workers.

Successful organizations also create a culture that values information as an asset and plans for information sharing, collaboration, and expert identification. In such an environment, business and IT

executives and managers view unstructured information as strategically as structured data and make use of that information in their daily analysis and decision making.

Unlike compute power, human expertise is not scalable. As AI systems take on the task of analyzing vast amounts of data, organizations must develop and empower knowledge workers, providing assisted support to these key roles. Knowledge workers will need to understand the relevance and applicability of different types of AI technology as they correlate AI and data with the processes of the business.

Combining both categories of preparation will ensure the team has the methods, discipline, skills, and resources to perform the invaluable task of extracting value from data.

An Eye on ROI

To get the maximum leverage out of these technologies, organizations must develop strategies, methods, and procedures for measuring success. Success metrics can also become key tools in the hands of top managers when promoting the use of information access and analysis solutions or when building a business case. Organizations should develop measurement strategies as part of their project and technology design and implementation plans. Areas of value or return can include outward-facing factors such as revenue growth, increased customer satisfaction, eliminated rework, increased productivity, and increased compliance. Internal value can come from increased employee satisfaction, reduced turnover, increased employee engagement, and faster ramp-up of new employees. Organizations that explicitly measure their information access and sharing efforts are much more likely to both gain significant benefits over those that don't and identify new opportunities for expanding the benefits that would not have been visible otherwise.

For example, when a leading manufacturer implemented information access and analysis technologies unifying file shares and other content management systems at its multiple facilities and offices, it saw improvements in productivity and realized a return on investment (ROI) based on increased innovation and faster and more accurate program bids. The company unlocked value from both historical and existing knowledge assets, which were unreachable prior to the implementation. The manufacturer has seen a general increase in the perceived importance of the unstructured information that the application is making accessible. Senior managers have been firmly behind both the project and the solution, and their own use of the application is growing.

Selecting the Right Project

A good adage is to start small and identify problems that have well-defined success criteria. The best initial projects have low decision risk, high repeatability, narrow data sets, and signal noise that can be mitigated using adaptable rules rather than manual investigations and still deliver positive results. Assess the data needed to ensure an accurate analysis and decision. Locate and license relevant third-party and public data, such as news or social media. Use AI to connect disparate data and then analyze linkages for potential answers.

Ensure cooperation within your enterprise and with its ecosystem, especially if your organization lacks experience with the new technology components. Building decision-centric iterative systems is a work in progress that requires collaboration among developers, domain experts, and data scientists. Innovative use cases not only can deliver significant business impact but also may address the perception of sustainable cost and scalability.

Fueled by Data

The availability of a wide variety of data and the technology, skills, and processes to take advantage of that data are radically changing how information is accessed, analyzed, and shared. Over 75% of respondents in an IDC survey believe that their organization already possesses data amenable for AI analysis and would benefit from external (third-party) data. The hidden value in content is unique to each organization, and unlocking it is not a trivial task. Many organizations face roadblocks because of a lack of appropriate technology and processes. However, organizations that do unlock this value have demonstrated an ability to increase revenue, improve productivity, reduce costs, respond to customer or stakeholder needs more quickly and accurately, and help bring products to market faster.

Many organizations are using internal websites, CRM and ERP systems, databases, email, instant messaging, and content management systems for storing current information, but large amounts of intellectual capital and value are also stored in documents across the organizations' file shares as well as legacy data stores. In addition, data comes increasingly from third parties, bringing the challenge of connecting data. These data sources hold the historical archive of knowledge assets that are not being used to their fullest potential, and organizations should include them as part of the information that needs to be accessed and shared.

Tools and Capabilities for Insights

Growth in this market continues to be driven by increases in artificial intelligence platforms and toolsets being used to develop applications ranging from chatbots and conversational interfaces to predictive and expert assist applications that offer advice and recommendations. AI software platforms include a plethora of technology components, such as:

- **Content analytics, information access, and search software** that analyzes, extracts, and organizes structured data from unstructured information and applies natural language processing to that consolidated data to help knowledge workers more quickly and accurately find, locate, and act on answers and pertinent information
- **Knowledge representation tools** such as knowledge graphs, triple stores, or other types of NoSQL data stores
- **Learning technologies** including machine learning, deep learning, and reinforcement learning, data science, and analytics (These platforms also provide for knowledge curation and continuous automatic learning based on tracking past experiences.)
- **Reasoning technologies** such as graph-based query languages or analogy-based query languages

Previously, early adopters of these technologies had to build custom applications. Today these technologies are being delivered as commercial enterprise applications and are often designed for replication or customization for additional areas and work patterns across the enterprise. Software capable of sending, receiving, and processing data in high volumes at high speeds is required. However, implementing streaming data integration tooling is more challenging than implementing the normal batch-oriented data integration tools. Data tools can quickly transform large amounts of raw data into consumable, quality information that is ready for analytics.

Cloud Environments and Security

It is widely known that the security capabilities offered by cloud service providers have much room for improvement. Because public cloud tenants generally lack the ability to build security in a cloud environment, it falls to cloud service providers to safeguard their cloud tenants' business security, and they are rising to the challenge. IDC is seeing a steady increase in the percentage of enterprises that

consider cloud service providers' security capabilities superior to their own organizations' capabilities. In other words, organizations remain concerned about security, but public cloud services are increasingly recognized as a more secure option than many self-run on-premise options. Providing customers greater assurance around their security concerns will remain a strategic differentiator for providers.

Security concerns and policy, legal, and compliance issues are very important to organizations, but those concerns need to be tempered with the attitude that information should be able to be shared with the widest number of people possible in the organization. Often, this requires a mental shift on the part of IT and policymakers to make information sharing the default and lock down or secure information only when necessary. Information access and analysis solutions should be able to tie into existing corporate security systems and verify that documents and other information have the right security profiles for appropriate users to access. In addition, organizations need to make that mental shift to share as much information as possible while maintaining their security standards. The default strategy should be to share information as widely as possible within the organization and with external stakeholders to promote collaboration, reuse, and innovation.

Personally Identifiable Information and Compliance Concerns

AI augmentation of knowledge work brings not only large potential benefits but also security and privacy issues about shared data, guidance, and recommendations that need to be handled appropriately within the enterprise. This has not been an inhibitor to adoption since the number of organizations using AI software platforms for fraud and compliance investigation increased significantly in 2016, especially within banking and finance.

CIOs and business leaders should prepare strategies to avoid and/or mitigate adverse impacts. Examples of components of such strategies include avoiding "black box" AI tools that defy understanding of how they work or reach decisions, working with stakeholders (employees, customers, partners, etc.) before deploying AI apps to understand their concerns and explain the benefits that they will gain, and engaging with a chief risk officer (CRO) to communicate the risks around AI adoption and develop corporatewide plans and strategies to mitigate these risks.

WATSON DISCOVERY

Portfolio and Components

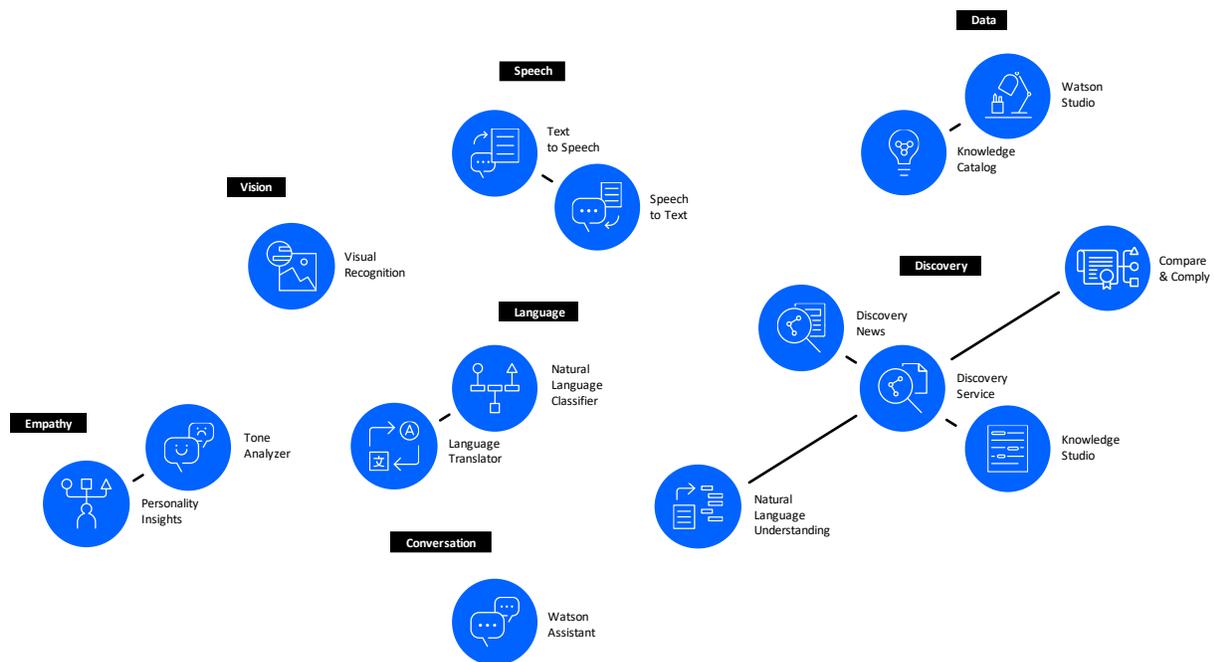
IBM has been an exemplar of technology leadership around data-driven analysis and decision making in technologies such as AI, machine learning, and other cognitive sciences. Part of IBM's Watson cloud-based AI platform, Watson Discovery Service incorporates or integrates with other Watson products such as Watson Natural Language Understanding (NLU), Watson Knowledge Studio, and Watson Assistant. Watson Discovery takes advantage of IBM Watson Studio as a foundation — providing a suite of tools for data scientists, data engineers, developers, and domain experts to build, train, and deploy custom machine learning and deep learning models at scale (see Figure 1). Watson Discovery users can quickly and easily create a custom insight engine with transformational capabilities such as:

- **The ability to perform natural language processing and analysis for both information ingestion and research to provide human-style communication (usually posed as questions and answers).** For example, with Watson Discovery's Element Classification contract analysis capability, users can discover the answers to questions such as "What are the obligations within my set of contracts that put my company at risk?"

- **The ability to statistically generate and evaluate series of evidence-based hypotheses to answer questions in a relevant and meaningful manner.** Discovery includes anomaly detection, which enables users to identify unusual data points, without additional coding, and take actions.
- **The ability to adapt and learn from training, interaction with humans, and outcomes related to the hypotheses generated.** Via relevancy training in Discovery, a machine learning model is trained to enhance result relevancy rather than force administrators to perform deep manual tuning of search engine parameters.
- **The ability to amass and enrich large amounts of structured and unstructured data needed to answer the questions posed.** Discovery provides automated ingestion options including tooling, API, and crawler technologies to bring data into the platform. Cross-collection federation allows users to search multiple collections with a single query. Passage retrieval enables users to find relevant information within documents and can extract snippets and return matching passages of text.

FIGURE 1

Watson Services



Source: IBM, 2018

Discovery is hosted on the IBM Cloud, allowing organizations to scale the environment to the needs of the application. Security and isolation are also handled by the IBM Cloud. As a cloud solution, the environment is managed and provisioned by the cloud service provider, allowing users to quickly begin working with Discovery's AI capabilities and features.

Guidance Through the AI Process

Discovery guides users through a process of ingesting data or content from one or multiple data sources, enriching that data with a broad selection of natural language processing capabilities (e.g., sentiment analysis, element classification, concept identification, and entity and relationship detection), and making discoveries or finding answers to questions by interrogating the corpus, that is, learning from and interacting with the results. IBM has implemented this as a continuous and iterative AI process, leading users to consider the last step first: What are the questions or hypotheses they want answered?

Start with Well-Prepared Data

At the heart of any AI insight engine is the data, which can consist of all manner of structured, semistructured, and unstructured data such as text, images, speech, and video. The Discovery user interface (UI) provides easy importing of data via APIs and has a wide variety of connectors to import data from many types of repositories (third party as well as users' own data). The world's news content comes out of the box in Discovery, with over 17 million news articles and blogs (both English and non-English) spanning the past 60 days and from thousands of sources.

Discovery aggregates data into a single enterprisewide "information fabric" that can be leveraged by a wide variety of analytics tools. Discovery starts with the documents in their native formats (such as PDFs and Microsoft Word documents) and then converts and normalizes that content so it can reside inside of Discovery. During this conversion process, Discovery applies natural language processing to add knowledge to the content by classifying it using configurable techniques such as sentiment analysis, document classification, and entity and relationship extraction. For example, Discovery can identify keywords and relationships between data so that the enriched data is prepared for subsequent analysis. Discovery's entity and relationship extraction can be customized to specific industries and domains (e.g., pharmaceutical versus retail) using Watson Knowledge Studio.

Enterprises have suffered through many and various problems that come from the amassing and storing of large data repositories, problems that stifle insights (i.e., the inability to find and extract relevant content) or cripple data governance. IBM's Watson Discovery consumes data with the intrinsic foundation of enabling users to easily and reliably find what they are looking for and then analyze the data to derive data-driven insights and drive actions.

Getting to Insights

After content is ingested and enriched, Discovery takes users through the query and insight stage. Discovery's administrative UI makes it easy to start building queries or searches via natural language or structured queries. With a schema explorer and visual query builder, users can easily build complex aggregate queries in a few clicks and also use key capabilities such as passage retrieval and document deduplication. With built-in machine learning for relevancy training, users can train Discovery to serve up the most relevant answers in their corpus, right in the administrative UI.

Because Discovery can ingest content continually, organizations can take actions and measure their effectiveness quickly and with the same queries. For example, negative sentiment on social media leads a municipality to change its bus route. The social media content is updated, the query is run again, and SMEs look for improvement in sentiment. This kind of measurement provides timely feedback on the effectiveness of business actions.

Availability

Watson Discovery is available on the IBM Cloud. IBM offers the Lite plan at no charge, allowing users to "try before you buy." Three paid tiers (standard, advanced, and premium) offer increasing levels of capabilities such as the number of documents and models, increasing customization, and security features. At each level, users can easily select and integrate other Watson services such as Watson Assistant and Watson Knowledge Studio. The latest pricing for Watson Discovery can be found on IBM's website or the catalog in IBM Cloud.

As of the time of the writing of this paper, several features are in beta, including knowledge graphs. Knowledge graphs in Discovery make connections across documents and generate new knowledge. Discovery can automatically create custom knowledge graphs from unstructured data by extracting and disambiguating entities and relationships, enriching the relationships using algorithmic techniques, and ranking the results using relevance algorithms. These knowledge graphs can function as the "knowledge hub" for organizations and can be applied for various use cases across enterprise search, recommendation engines, and other decision-making processes — for example, detecting fraud, waste, or abuse. The use of a custom model (created in Watson Knowledge Studio) in the knowledge graph creation process can help build domain-specific knowledge graphs with applicability in domains such as financial, technology, security, intelligence, and healthcare.

Client Examples

Australian data analytics and software engineering agency Max Kelsen chose Watson Discovery Service to bring the voice of the customer to a local government entity. "The thing that leads us into Discovery and Watson generally is we saw a mature suite that we were able to build on top of and get things up running a lot faster than we would if we did it ourselves," said CEO Nick Therkelsen-Terry. Discovery allowed the company to create a custom insight engine and deliver to its customer within 12 weeks instead of 12 months using another platform. Discovery's toolset and easy-to-use interface allowed the company and its analysts to focus on the critical issues (questions) for the company's client. They were able to quickly move from broad sentiment to specific issues to design a response or advise the client on actions to improve customer experience. Watson Discovery provided the natural language processing combined with the ability to put it into a searchable index that has the metadata enriched for it and a rich query interface that allowed the company to create analytical queries that it would not be able to create with a straight search engine. The time from data setup to query output was reduced from two weeks to five minutes. In addition, the models the company created were adaptable from one corpus of data to others. The company found the Watson Discovery tools far ahead of the competition, citing capabilities such as the combination of machine learning, text analytics, text extractions, metadata enrichment (combined with indexing of searching capabilities and putting the index into a knowledge graph), and a good query interface.

Extending the project to new clients, the company is seeing interest from financial institutions, government, and healthcare, offering its clients actionable insights on a reasonable budget (both cost and time). Fast time to actions and real-time tracking of results allow its clients to rapidly reiterate and tune their actions to achieve the desired business results.

Oovvuu, an AI start-up in Australia, created new revenue and new clients by delivering enriched media content in breaking news. Compass, Oovvuu's proprietary AI application, uses Watson Discovery News and Natural Language Understanding to ingest and enrich hundreds of thousands of news articles and videos. It quickly analyzes, categorizes, and recommends videos to be placed alongside web-based news stories, resulting in maximum reader/viewer exposure and increased web page engagement, which

drives increased advertising revenue for media and news organizations. The use of Watson Discovery News allows Oovvuu to process and add video to stories from 143 countries and 100,000 publishers at a rate of 300,000 per hour. The potential revenue over five years is US\$20 billion.

Oovvuu needed immediate real-time categorization combining the huge volumes of videos and news stories across the globe and a cloud-based service that scales with the organization provided by Watson knowledge discovery services. IDC named Oovvuu a Digital Disruptor of the Year in 2017, a direct result of the company's work with IBM Watson. Oovvuu plans to also offer image recognition and video scene segmentation using Watson services in the future.

CHALLENGES/OPPORTUNITIES

But IBM does face challenges. First and foremost, Watson Discovery is a new technology approach that either complements or replaces existing approaches. As such, IBM must commit time and resources to educating target customers about the benefits of implementing the cognitive system. (Results will vary, but IBM is targeting deployment times of three to six weeks and a six-month ROI.)

The market views Watson either as a large computing system requiring a significant up-front investment or as a marketing effort. Over the past two years, IBM has worked diligently to move more and more Watson capabilities to the cloud, making up-front investment minimal. As for the latter, IBM has eliminated the tendency to attach Watson to every new product and has reinforced the brand by bringing innovative products such as Watson Discovery to market.

However, to be successful with AI-enabled applications such as Watson Discovery, IBM must identify and promote use cases where the abilities of the tool to make connections and inferences result in significant cost savings and return on investment. Developing playbooks for common Watson Discovery use cases will help solidify the stories and opportunities where Watson Discovery Service can be best applied.

CONCLUSION

Today, organizations are rapidly falling behind in effectively using the full amount of information they collect and store. Even more important, they're not locating the information they need to make decisions that affect their organizations daily. In addition, competition and disruption with digital transformation using AI are becoming common. Organizations need to take the first step to address this set of needs.

The modern worker must be adaptable, and increasingly AI technology will be a key enabler of this adaptability. However, the biggest shift is in the activities themselves. The modern knowledge worker applies experience and knowledge to business decision making using new collaborative tools, data/analytics and, increasingly, technologies such as AI-enabled applications (i.e., AI insight engines). The ability to synthesize inputs (data, predictive, and prescriptive analytics and recommendation systems and assistance) into business decisions and strategies to drive better outcomes is the new nature of work. A wide variety of solutions and tools is being brought to the market that provides this level of augmentation for knowledge workers, provided by a range of traditional and new vendors, such as IBM and its partners. Organizations need to understand that these systems will change the very nature of work in their organizations and provide a level of empowerment to knowledge workers based on evidence and data that few have ever had available before.

IDC believes that AI-enabled applications such as AI insight engines will become ubiquitous and that insight engines using AI and machine learning capabilities will have a significant impact on the workplace. Because these systems are essentially power tools for the human mind, the use of these tools will become a requirement in the future. Almost every discipline that uses or touches information in one way or another is a candidate for these next generations of systems. Many organizations are already evaluating or using these technologies as mainstream tools in day-to-day operations.

Enterprises should consider using assistive technologies including tools such as AI insight engines. Prototypes and pilot projects using these tools are a great way for staff to become acquainted with the power and capabilities that the tools deliver so that, over time, more widespread use can drive significant competitive advantage and operational improvements.

Specific recommendations for enterprises include the following:

- **Match business outcomes to the capabilities of technology offerings.** Organizations should make sure the AI solution that is chosen will be able to help achieve the desired business outcome and/or address the issues that are planned to overcome utilizing an AI system such as an AI insight engine.
- **Develop an information strategy and build the knowledge base.** To build a sufficiently capable knowledge base, organizations must collect all pertinent and domain-related information sources, including unstructured data sources. Organizations that have strong content and knowledge management resources will be able to do this more easily than organizations that do not have such resources.
- **Identify key stakeholders and specific functions within enterprise content or data management that would benefit from the use of advanced tools such as an AI insight engine.** For example, these may include meeting compliance requirements, litigation discovery needs, competitive analysis, research and development, or other types of business analysis. Focus first on business processes that have the most to gain from leveraging AI technologies — that is, greatest cost savings, enhanced risk mitigation, increased automation, or highest value added.
- **Consider the use of an AI insight engine such as Watson Discovery Service to provide automated analysis and knowledge discovery services that will demonstrably improve the organization's bottom line.**

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2018 IDC. Reproduction without written permission is completely forbidden.

