

The Messy World of the Modern Data Spike

If data were visible, we would see nothing else. Every day, we create over 15 Petabytes (1 PB = 1,000 Terabytes or 1,000,000,000,000,000 bytes) of new information.ⁱ Every minute, 100 or more LinkedIn accounts are created.ⁱⁱ Every second, one hour of video is uploaded to YouTube.ⁱⁱⁱ

This spike of data comes with tremendous opportunities. Experts around the world are producing valuable new information for those who can find it. About 2.5 million scientific articles were published in 2014, in one of over 28,100 active scholarly peer-reviewed English language journals.^{iv} Interested in patents? American companies filed 57,385 patents in 2015, ahead of Japan (44,235) and China (29,846).^v

This constant stream of new information can provide great value to businesses looking for new insights. But we're far past the point where humans can keep up. Valuable information remains buried under the data avalanche. People need help.

Technology, in the form of programmable systems, plays a critical role in organizing data. But in the past, analyzing data has required people to make sense of the information. People were needed to find value in the data by reading emotional cues, understanding nuance, making logical assumptions, and making recommendations. Today, cognitive computing applications are helping people and companies all over the world gain new insight from their data. Developers, startups and enterprises are expanding the capabilities of their applications and reaching new markets by incorporating cognitive technology into their own applications. The trick? Choosing the right cognitive computing tools and partners.

Cognitive Computing Systems

IBM defines cognitive computing as “systems that learn at scale, reason with purpose, and interact with humans naturally.” They go far beyond simple tabulation and calculation to perform functions that previously could only be accomplished by humans. Things like: recognizing images, understanding spoken and written word, and knowing when to apply information to different situations. And just as humans learn over time, cognitive technology continues to “learn” as well. Cognitive technology is truly ushering in a new era of computing, where intelligent systems will aid us in making better decisions across all domains.

In order to support the growth of cognitive applications, IBM has made cognitive technology available to developers through the Watson Developer Cloud. The APIs available on the Watson Developer Cloud enable developers to build cognitive capabilities into their existing applications or create brand new products with a cognitive core.

Drawing Insights from Unstructured Text

As analyst firm IDC reports, “unstructured content accounts for 90 percent of all digital information.”^{vi} Unlike the structured data you find in excel documents and databases, unstructured data, particularly unstructured text, is everywhere. And as humans, we interact with unstructured text every single day - in social media postings, the emails we read, company reports, news articles, information on websites, and so on.

The numbers for email alone are staggering: nearly 2.6 billion email users today send over 200 billion messages daily.^{vii} Those numbers will continue to increase to 2.9 billion users and over 246 billion messages in 2019.^{viii} Add that to the flood 500 million or more tweets per day.^{ix} Every 60 seconds, Facebook’s one billion or more active daily users^x publish over 510 comments, 293,000 status updates, and upload over 130,000 photos. The flood of information from social media data has become overwhelming.

Until very recently, computing systems struggled with analyzing, gaining insights and making use of this vast spread of information. But help is now available. IBM Watson is enabling startups and developers to begin making use of this data through the Watson Developer Cloud. Developers can now tap into a wide range of cognitive technology, which understands natural language.

Watson API's at Work

Analyzing Text

One service available on the Watson Developer Cloud, AlchemyLanguage, provides text analysis functions through natural language processing. The AlchemyLanguage APIs can analyze text and help users understand sentiment, keywords, entities, high-level concepts and more.

How can this service be used? Pulsar^{xi} is a social intelligence platform focused on social media monitoring that uses the AlchemyLanguage API to provide brands with valuable customer intelligence. To help their clients stay current, Pulsar worked with IBM Watson to read and analyze between 20 and 30 million documents per month across all major social media platforms. Pulsar’s platform goes far deeper than keywords, providing clients with key insights so they can position themselves around current trends, discussions, and topics. For example, Pulsar can tell an automotive manufacturer not only what cars their customers are talking about, but also which design features they care about most. Through the integration of several natural language processing APIs, Pulsar rapidly releases innovations that help customers go beyond basic keyword tracking to map brand audiences, observe how content travels and manage team communication.

Another popular API, The IBM Watson Dialog service, powers comprehensive and robust conversations between virtual agents and users. The Dialog service enables applications to use natural language to automatically respond to user questions, cross-sell and up-sell, walk users through processes or applications, or even hand-hold users through difficult tasks. The Dialog service can track and store user profile information to learn more about end users, guide them through processes based on their unique situation, or pass their information to a back-end system to help them take action and get the help they need.

How can this service be used? One company used the Dialog service to create a cognitive maintenance and repair application for the oil and gas industry. Maintenance technicians, tasked with keeping up with evolving environmental legislation, can directly query legislation to get the answers they need saving time and money.

Making Use of Audio

Many only consider text when evaluating unstructured data. But in this multi-media world, audio has become just as important. From podcasts, to videos, to voice-controlled applications, there is a growing need for tools to help us analyze audio data. In 2013, Apple said iTunes had over 1 billion podcast subscribers to over 250,000 unique series in over 100 languages.^{xii} These numbers exclude the hundreds of thousands of other podcasts on distributors like SoundCloud.com, Libsyn.com, Radiotopia, Infinite Guest, and Soundworks. The IBM Watson Developer Cloud provides developers with tools to translate audio information.

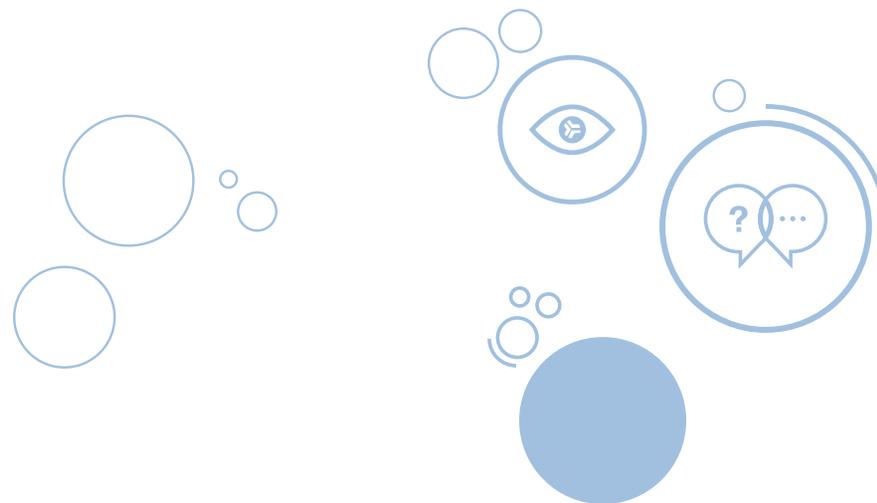
Watson's audio capabilities, like the Speech to Text API, help developers with tasks like transcribing meetings and enabling voice controlled applications. They can also be used in combination with other APIs to accomplish more complicated tasks like creating personality profiles and analyzing call center recordings to understand customer sentiment.

Watson API's at Work

Analyzing Audio

The Speech to Text API can be used anywhere there is a need to bridge the gap between spoken word and its written form. This easy-to-use service uses machine intelligence to combine information about grammar and language structure with knowledge of the composition of an audio signal to generate an accurate transcription. It uses IBM's speech recognition capabilities to convert speech in multiple languages into text for a wide range of purposes.

How can this service be used? The Speech to Text API has a wide range of applications. Several developers have used the Speech to Text service to transcribe media files, enable voice control applications, and enable note dictation.



Intelligently Identifying Images and Video

Image and video recognition provides another tremendous opportunity for cognitive computing to power innovation. Traditionally, only people could understand and process the information found in images and videos. Images needed to be tagged by people so that others could find them. With advancements in image recognition capabilities, many companies are starting to take advantage of these growing opportunities.

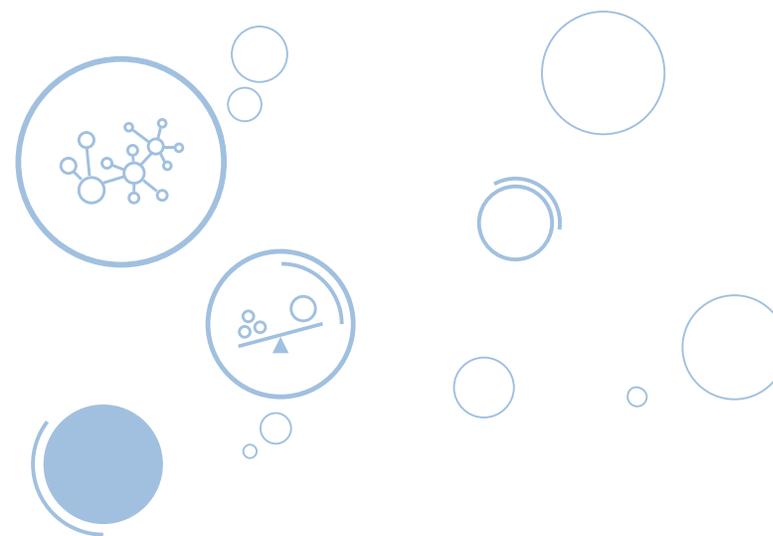
Image recognition capabilities available in the IBM Watson Developer Cloud can be used across a wide variety of domains. Marketers are finding that they can provide highly targeted marketing offers based off of images that users are uploading to popular social media sites. Healthcare providers are looking to use image recognition technology to assist doctors in analyzing X-rays and other medical imagery. And many consumer applications are starting to pop up, helping users find retail products they see in pictures.

Watson API's at Work

Analyzing Images

The AlchemyVision API employs deep learning innovations to understand a picture's content and context. The API can analyze images and return the objects, people, and text found within the image.

How can this service be used? One company, Tabelog, uses AlchemyVision to enable foodies to discover the best dining experiences across the US. Tabelog pulls images from social media and analyzes them along with restaurant reviews to provide users with appropriate expectations for local restaurants.



Building Next Generation Cognitive Applications with IBM Watson

The IBM Watson Developer Cloud provides a library of cognitive tools to help developers bring cognitive computing to their applications. Many startups have brought new innovations to market that leverage the value and technical foundation of IBM Watson. Some are fun, like making better fantasy football picks and streamlining vacation travel plans. Others are serious, like personalized health and wellness tools. Others are life-saving, like reading radiological images to spot anomalies. Here are just a few examples:

Asklepius, created by Best Doctors, brings centralized medical expertise to users in real time. This application offers targeted, real-time answers to users' queries about their medical health. 95 percent of Best Doctors members report high satisfaction ratings.

Genie MD utilizes Watson to deliver personalized health suggestions. Their mobile Patient-Centered Care Coordination Platforms enables patients and their caregivers to keep track of their health profiles across multiple devices, as well as easily diagnose potential issues.

@Point of Care is a medical education advisor that offers clinical decision support. Their Watson-powered solution allows clinicians to access curated, evidence-based, peer-reviewed, disease-specific content on a mobile platform at the point of care.

50 Wise removes complexity from the payment world and all services related to credit cards. The value of their eVision application comes from leveraging Language Identification and Machine Translation to automate all credit card services for the cardholder.

Red Ant closes the retail clerk training gap with intelligent sales support. Their Sell Smart mobile app uses Watson to analyze a vast corpus of product information and technical specification to provide product recommendations. Watson Personality Insights helps tailor the service to specific customers, increasing the effectiveness of the salesperson.

AdvisorQA is a client-facing social collaboration tool for investors and advisors. Built for Financial Advisors in the Wealth Management Industry, it primarily targets Millennials who are mobile and interested in investing. Clients using AdvisorQA can easily make sense of market sentiment derived from news articles and social media, and discover interesting trading opportunities.

Sellpoints provides insights to retailers about online customer searches. Their Natural Selection application, powered by Watson, understands a shopper's intent through analyzing natural language and guides consumers to the right product in as few as two simple clicks.

Many other startups use Watson to bring cognitive computing to their applications. The Watson Developer Cloud has vibrant community of thousands of developers in over 85 countries.

Getting Started with Watson

Whether you have an existing application, or would like to start from scratch, the Watson Developer Cloud offers several APIs to help you bring cognitive technology to your application. The real power of IBM Watson comes from the ability to mix and match APIs in a wide variety of ways. To help jump start your development process, IBM has put together 3 application starter kits. These starter kits explore some of the popular ways that Watson technology is being used and provide code examples to get you up and running quickly.

Business Intelligence

What can you do?

Analyze news and social media to understand how people are talking about your company or learn about important topics and how people feel about these topics.

Services Used

Alchemy Language + AlchemyData News + Tone Analyzer

Example Application:

Business Intelligence Application – This sample application demonstrates how to use web and social media data along with natural language processing techniques to understand how people feel about companies.

[Launch App](#)

[Get Started with the Business Intelligence Application Starter Kit](#)

Audio Analysis

What can you do?

Analyze news and social media to understand how people are talking about your company or learn about important topics and how people feel about these topics.

Services Used

Alchemy Language + AlchemyData News + Tone Analyzer

Example Application:

Business Intelligence Application – This sample application demonstrates how to use web and social media data along with natural language processing techniques to understand how people feel about companies.

[Launch App](#)

[Get Started with the Business Intelligence Application Starter Kit](#)

Conversational Agent

What can you do?

Give your users an engaging experience by enabling them to interact with an automated conversational agent.

Services Used

Dialog + Natural Language Classifier

Example Application:

Movies Conversational Agent – This sample application allows you to chat with Watson to discover new and upcoming movies based on genre and rating.

[Launch App](#)

[Get Started with the Conversational Agent Application Starter Kit](#)

Discover the full suite of APIs available on the Watson Developer Cloud and start your free 30-day trial today. For more information visit: ibm.com/WatsonDeveloperCloud

References

ⁱ<http://www.slideshare.net/adlibsoftware/presentation-ecm-and-dark-data>

ⁱⁱ<https://www.linkedin.com/pulse/hr-data-big-yet-arun-krishnan>

ⁱⁱⁱ<https://www.youtube.com/yt/press/statistics.html>

^{iv}http://www.stm-assoc.org/2015_02_20_STM_Report_2015.pdf

^v<http://phys.org/news/2016-03-world-leader-patent.html>

^{vi}https://idc-community.com/groups/it_agenda/bigdataanalytics/unlocking_the_hidden_value_of_information

^{vii}<http://www.radicati.com/wp/wp-content/uploads/2015/02/Email-Statistics-Report-2015-2019-Executive-Summary.pdf>

^{viii} *ibid*

^{ix}<http://www.internetlivestats.com/twitter-statistics/>

^x<http://newsroom.fb.com/company-info/>

^{xi}<http://www.pulsarplatform.com/>

^{xii}<http://www.macworld.com/article/2044958/apple-one-billion-itunes-podcast-subscriptions-and-counting.html>