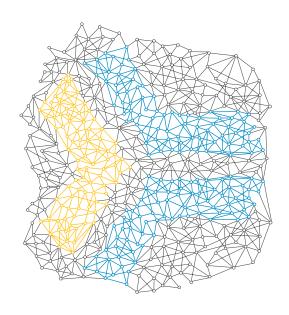


Deep Learning: 6 Real World Use Cases



Practical Applications for Today's Data-Driven Business

Written By Janet Wagner for AlchemyAPI





Solving Big Data Problems With Smart Machines

Deep learning is everywhere. It's on Amazon and Netflix making personalized recommendations. It's on your smartphone helping your voice-activated assistant understand you. It's helping websites and mobile applications transform content into precisely targeted advertising. It's helping companies gain meaningful insights from unstructured data.

The vast majority of data generated from the internet is unstructured; it is data intended for human consumption and not designed for computers to process. In a session at GigaOM Structure Data 2014, Elliot Turner, CEO of AlchemyAPI and Stephen Gold of IBM Software Group, Watson Solutions, spoke about unstructured data and deep learning. During their session, they discussed how smart machines are enabling and empowering today's modern organizations.

"We're seeing companies leverage unstructured data—things like photographs, videos, chat logs, documents—to make better, more informed business decisions to automate processes. They're leveraging human-like capabilities inside automated workflows with deep learning. I think that these technologies can ultimately augment what's possible in business and humanity, but not necessarily replace it," shared Turner.

Gold added, "The vast form of data that's available to us is all unstructured. It is the text, the blogs, the tweets, the articles, and the photographs – and we're just now starting to experience this idea of a system that is readily approachable that you can navigate through natural language; not through speech recognition, but actually through understanding where it literally can read and comprehend in context the data itself."



"I think that these technologies can ultimately augment what's possible in business."

Elliot Turner, CEO at AlchemyAPI





Big Data? This Is Massive Data

According to an IBM study, approximately 2.5 quintillion bytes of data are created every day and the IDC is estimating that by 2020, approximately 40 zettabytes (ZB) of data will be generated every day. The Internet Live Stats site (part of the Real Time Statistics Project), provides real time updates showing just how quickly data is being generated on the internet.



3,024,476,315

Internet Users in the world



1,139,737,171

Total number of Websites



145,692,209,447

Emails sent today

g

2,776,526,173

Google searches today



2,547,826

Blog posts written today



486,695,263

Tweets sent today



5,516,365,773

Videos viewed today on YouTube



92,250,973

Photos uploaded today on Instagram



97,066,147

Tumblr posts today

http://www.internetlivestats.com Screenshot: 12/9/14

With numbers like these, it is clear to see how so much value can be contained, and sometimes hidden, in all the data collected by a business. Not only is it important to recognize where the data is held, but it is paramount to extract real insights to understand consumers and make better decisions.

Without the full story behind the data, organizations play an expensive guessing game as to what consumers need and want. The good news is that deep learning systems give organizations the ability to capitalize on the data at their fingertips and ensure they stay agile, focused and knowledgeable.





What is Deep Learning?

Development teams are overwhelmed with requests to create applications that automatically gather and synthesize data so that organizations can make better content and purchasing recommendations, extract keywords for search engine optimization (SEO), collect brand intelligence to develop effective messages and more. Many teams support their data analysis efforts with deep learning to remain innovative, speed time-to-market, and boost revenue.

In a recent AlchemyAPI webinar, Aaron Chavez, Chief Scientist at AlchemyAPI, explains the idea of deep learning saying that "the general idea of deep learning is to use neural networks to build multiple layers of abstraction to solve a complex semantic problem."

In other words, deep learning is a new area of machine learning that works to improve things like computer vision and natural language processing to solve unstructured data challenges.

Read on to learn six practical applications for solving real-world problems with deep learning today.

The general idea
of deep learning is to
use neural networks to
build multiple layers
of abstraction to solve
a complex semantic
problem."

Aaron Chavez, Chief Scientist at AlchemyAPI

What is Deep Learning?

5 Deep Learning Resources Everyone Should Bookmark

Learn More





Voice Search/Voice-Activated Assistants

One of the most well known and popular uses of deep learning APIs is to power voice-activated intelligent assistants, a feature found on nearly all smartphones. Most consumers are familiar with the major players in the smartphone OS market: Apple, Google, and Microsoft. Apple's Siri is the most well known of the voice-activated assistants and has been around the longest (Siri was introduced in October 2011). Google Now, the voice-activated assistant for Android, was launched less than a year after Siri. The newest of the voice-activated intelligent assistants is Microsoft Cortana which was introduced in April of this year on Windows Phone 8.1.



In November 2013, Microsoft introduced natural-language voice search on the Xbox One console. Using Xbox One's voice navigation users can search for movies, music, and other content using natural phrases. A few months ago, Microsoft introduced Azure Media Indexer, a new Azure cloud platform service that allows users to index video content based on the words spoken in the video. The service uses natural language processing (NLP) technology, which is based on deep learning, to make media files and content searchable and to augment the platform's vocabulary.

APIs that allow developers to add speech recognition functionality to their applications include the AT&T Speech API and the API.AI API.



Recommendation Engines

Recommendation systems have become a very popular and common feature of web and mobile applications. Major companies such as Netflix, Amazon, Google, Facebook, and Twitter have access to a vast wealth of user generated data. This access to data has allowed these companies to implement complex recommendation systems that provide added value to both users and the companies themselves.

According to a recent SlideShare presentation by Xavier Amatriain, research/engineering manager at Netflix, 2/3 of the movies users watch on Netflix are recommended, Google News recommendations generate 38% more clickthroughs, and 35% of Amazon's sales are generated from recommendations.

Collaborative filtering and content-based filtering have traditionally been the two common approaches used in the design of recommendation systems. In recent years recommendation systems have become far more intelligent with companies using deep learning to predict user preferences and provide accurate recommendations. A recent Netflix blog post explains how the company is researching deep learning and other machine learning techniques to improve recommendations and provide greater personalization to users. A Spotify intern is currently working on a deep learning approach to music recommendations called "deep content-based music recommendation," described in this blog post.

The winner of last year's MHacks hackathon (AlchemyAPI was a sponsor) was Assimilator, a content recommendation engine that "automatically finds related articles and inserts them into product pages on ecommerce sites." The application uses AlchemyLanguage APIs to scan web pages for keywords and then uses those keywords to find related content.

facebook.













Image Recognition

The goal of image recognition technologies is to recognize and identify objects in images as well as understand the content and context. Google and AlchemyAPI have been developing image recognition technologies for quite some time. In June 2012, Google brought much attention to deep learning through the publication of a paper highlighting the company's image recognition research.



AlchemyVision's Face Detection and Recognition service is able to distinguish between look-alikes such as actor Will Ferrell and Red Hot Chili Peppers' drummer, Chad Smith.

The paper details how Google built a network and models capable of identifying cats and human faces using unsupervised learning. The network was also able to recognize cat faces and human bodies. A few months ago, AlchemyAPI announced the public release of the AlchemyVision Face Detection and Recognition API providing developers programmatic access to the company's image recognition technology.

CamFind is a mobile application that recognizes and identifies objects in photos taken by mobile phone users. CamFind utilizes AlchemyVision APIs which allow the app to tell users what objects are in photos and is even capable of telling users where items can be purchased. One of the key features of AlchemyVision is the ability to recognize multiple concepts within the same image.



Image Tagging/Image Search

Image tagging APIs provide applications with the ability to recognize and categorize many different types of objects in images including animals, buildings, nature scenes, people, faces, etc. Many image tagging APIs are also able to recognize and understand different concepts within the images. They also provide applications to take an image and then search for similar images based on the original.

Google uses image tagging technology to allow Google+ users to search their photos by content without having to tag the photos beforehand. Facebook is using image tagging to improve the photo sharing experiences of users.

SimplyMeasured, a leading social media marketing and analytics company was an early tester and proponent for computer vision APIs, using them to tag, classify, track and measure the images used in campaigns which better targets customers and boosts sales. CrisisNET, an Ushahidi initiative described on the official website as "the firehose of global crisis data," uses image tagging and other AlchemyVision APIs to tag, classify, and analyze images pulled in from thousands of government, business, humanitarian, and other data sources.

Learn how CrisisNET uses AlchemyVision to gather crucial information for fast-acting news organizations.

Read the Story

As shoppers become increasingly comfortable making their purchases online, major brands are driving a large amount of sales through eCommerce offering the ability to track and measure everything in a campaign, including the images used, gives brands a competitive advantage when targeting customers and driving sales. With AlchemyVision, we have been able to accurately tag and classify a good portion of images at very high rates with minimal human effort."

Aviel Ginzburg, CPO at Simply Measured









More and more companies are turning to deep learning when it comes to marketing and advertising brands, products, and services. Deep learning makes it possible for ad networks and publishers to leverage their content to create data-driven predictive advertising, real-time bidding (RTB) advertising, precisely targeted display advertising, and more. Baidu, the Chinese version of Google, is using deep learning to precisely predict advertising that is relevant to users which has helped significantly increase the company's revenue.

Earlier this year, AlchemyAPI launched new deep learning-based taxonomy and sentiment APIs designed to transform content into valuable digital ad inventory. The Taxonomy API is "IAB-compliant and includes an expanded taxonomy of over 1000 topic categories with up to five levels of hierarchy. The self-service component of the API makes it possible to create categories when given any arbitrary phrase."

Companies that are leveraging the Taxonomy API include (RTB) mobile advertising platform AdTheorent and the industry prospecting productivity tool FlockMiner. AdTheorent uses AlchemyAPI keyword extraction and sentiment analysis APIs to power their predictive modeling engine which has helped to significantly improve click-through rates. FlockMiner uses the Taxonomy API to bridge gaps in the platform's people suggest tool.



Pattern Recognition

Pattern Recognition allows companies to monitor and process a multitude of things. Here are two very different use cases:

PayPal

PayPal is using deep learning via H2O, an open source predictive analytics platform, to help prevent fraudulent purchases and payment transactions.

H2O uses advanced machine learning algorithms to analyze data for anomalies indicating fraudulent activity and security threats in real time.



Enlitic

New startup Enlitic is using deep learning to process X-rays, MRIs, and other medical images to help doctors diagnose and treat complicated diseases.

Enlitic uses deep learning algorithms that "are suited to discovering the subtle patterns that characterize disease profiles."





Conclusion

These are just some of the real-world use cases for deep learning. There are many others such as speech translation, speech transcription, text classification, content discovery, brand intelligence, and social media monitoring.

Deep learning is a disruptive technology that is being used by more and more companies to create new business models and to build innovative applications that help solve real problems. While there may not be one industry-accepted definition of what deep learning is, there is an industry-accepted agreement as to the benefits and exciting applications deep learning has made possible, and will make possible in the future.

Get Started With AlchemyAPI

Relieve yourself of the cost and complexity of building your own system.

Learn More >

Read Our Customer Success Stories >

Register for a Free API Account >

Contact Us

Contact Sales

888-777-7405

alchemyapi.com

@AlchemyAPI

About AlchemyAPI

AlchemyAPI's mission is to power smart applications that understand human language and vision by making breakthroughs in deep learning-based artificial intelligence available to everyone. AlchemyAPI is used by more than 40,000 developers across 36 countries and a wide variety of industries to process over 3 billion texts and images every month. For more information, visit our website at alchemyapi.com.

About Janet Wagner

Janet Wagner is a data journalist and full stack developer based in Toledo, Ohio. Her focus revolves around APIs, data visualization, machine learning, and data-driven journalism.

Follow her on Twitter: @webcodepro, Google+, or send her an email.

© 2014 AlchemyAPI