Using Deep Learning for Sentiment Analysis and Opinion Mining

Gauging opinions is faster and more accurate.

Abstract

How does a computer analyze sentiment? How does a computer determine if a comment or a personal review is positive or negative? Sarcastic? Biased?

The word *sentiment* refers to an attitude, feeling, or emotion associated with a situation, event, or thing—an opinion—which can be difficult to quantify, even using traditional modes of opinion mining or sentiment analysis.

Deep learning and inference dramatically improve sentiment analysis in two ways:

- Increasing accuracy
- Making opinion mining much more useful

Gauging Response

A customer comments on a product—or a restaurant, say, or a movie or a political candidate—on social media. If the comment says "Excellent!" the sentiment of the review is easy to determine (and a human reading that comment would make the right determination of its sentiment every time). But how does a computer analyze sentiment? How does a computer determine if a comment or a personal review is positive or negative? Sarcastic? Biased? And how can a computer effectively analyze sentiment—the feelings behind words—on a grander scale (say, the Twittersphere)?

The technology of Sentiment Analysis, or opinion mining, enables users to understand how large groups of people are responding to a product, service, company, or idea.

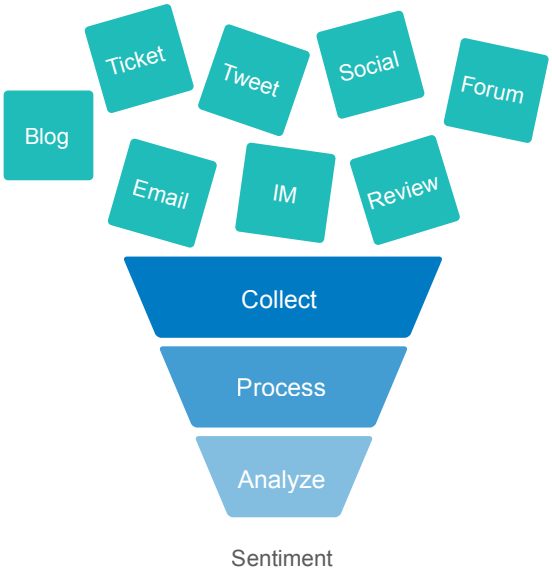

Figure 1: The basic flow of the sentiment analysis process: collect (using a wide variety of inputs, including social media, blogs, tweets, email, and others); process (to look for words that indicate sentiment), and analyze (parse to refine and weight sentiment).

Quality and engagement metrics generally lack enough context to be truly accurate or useful. Sentiment analysis adds context to them.

Many companies are already using all kinds of data—freeform comments, social media shares and likes, re-tweets, inbound links, and onsite engagement—as invaluable metrics that demonstrate how people respond to products or services (see Figure 1). These metrics are excellent at measuring engagement. But they are not as useful at measuring how people feel about a product or service.

Quality and engagement metrics generally lack enough context to be truly accurate or useful. Sentiment analysis adds context to them. Aggregated and contextualized sentiments can indicate trends and patterns to provide a much better and more accurate picture of brand reputation.

Old School: Rules-Based or Corpus-Based

Traditionally, sentiment analysis systems have looked at words in isolation. Typically, they assign positive points for positive words and negative points for negative words, then sum up these points. In the most simple approach, rules-based sentiment analysis, phrases are broken down into parts-of-speech trees. For example:

"I love this microwave!"
The word "love" generates a positive "+1" ranking.

"This is a terrible airline and I hate the customer service."
The words "terrible" and "hate" generate a -2 ranking.

Words like "love," "terrible," and "hate" would be stored in a "sentiment dictionary" to detect positive and negative words. These dictionaries can also detect "booster" word techniques. For example:

"The coffee was really, really bad."
"Bad" by itself would have earned a -1 ranking, but the "really, really" booster generates a -2 ranking.

Various forms of advanced linguistic analysis can be performed on the text to detect features like negation:

"I did not find any problem with the service, and the food was great."
This gets a +2 ranking. Even though the word "problem" appears in the comment, it's negated by the phrase "did not."

This rules-based form of classical sentiment analysis requires manually crafting the text analysis and parsing rules, efforts that are time-consuming and labor-intensive. This approach is also harder to transfer to other languages, and does not work well with channels like Twitter which have shorter, condensed, and idiosyncratic sentences.

Another approach to traditional sentiment analysis is called "corpus-based" (see Figure 2); it uses a double propagation

This rules-based form of classical sentiment analysis requires manually crafting the text analysis and parsing rules, efforts that are time-consuming and labor-intensive.

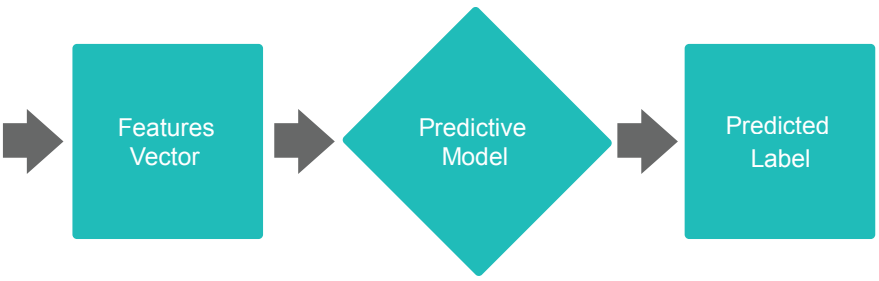

Figure 2: In a corpus-based sentiment analysis model, a features vector and a predictive model generate a predicted label: Is a sentiment good or bad?

between opinion words and the items they modify. It requires a large corpus to get good coverage and depends on exhaustive feature modeling. The accuracy of the analysis depends largely on the types of channels, their length, and the corpus models available.

Accuracy of the best sentiment analysis tools available with these two classical methods is approximately between 40% and 60%—which is good but not outstanding.

New School: Applying Deep Learning

To be useful, sentiment analysis must be accurate—and by applying deep learning to the challenge, many organizations are making radical leaps forward in terms of accuracy and utility. Recursive Neural Networks (RNN) using the Long Short-Term Memory (LSTM) architecture on top of grammatical structures can help provide more accurate measurement of sentiments in both small and large bodies of texts across different channels (see Figure 3). Sentiment analysis using deep learning does not need handcrafted features or a rigorously determined dictionary; instead, it uses inference to create its own models.

To be useful, sentiment analysis must be accurate—and by applying deep learning to the challenge, many organizations are making radical leaps forward in terms of accuracy and utility.

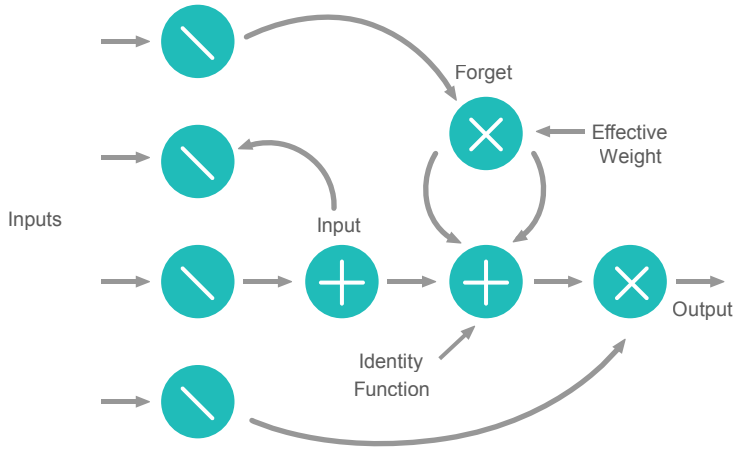

Figure 3: Deep learning sentiment analysis using RNNs uses LSTM-based memory cells, which can analyze/identify an input, "forget" the input if necessary (i.e., when it's irrelevant), and weight the result appropriately. It retains memory between training examples to make sentiment analysis more valuable.

RNNs have connections with loops, adding feedback and memory to the networks over time. This allows an RNN to learn and generalize across sequences of inputs as opposed to just individual patterns. They are very similar to complete programs and are "Turing complete."

RNNs can predict positive or negative sentiments as a binary classification task. In a traditional deep neural network, during the gradient back-propagation phase, the gradient signal can end up being multiplied a large number of times by the weight matrix associated with the connections between the neurons of the recurrent hidden layer. This means that the magnitude of

weights in the transition matrix can have a strong impact on the learning process.

LSTM is a powerful type of RNN. It introduces a new structure called a "memory cell," which includes an input gate, a neuron with a self-recurrent connection, a forget gate, and an output gate. With their ability to retain memory between training examples, RNNs allow capture of relations between words and provide the ability to remember important information across longer durations of time.

Faster and More Accurate

Using deep learning, sentiment analysis is much more effective than classical methods, dramatically improving both speed and accuracy. As noted, traditional sentiment analysis delivers accuracy of between 40% and 60%, but sentiment analysis powered by deep learning can achieve accuracy of between 80% and 90%.

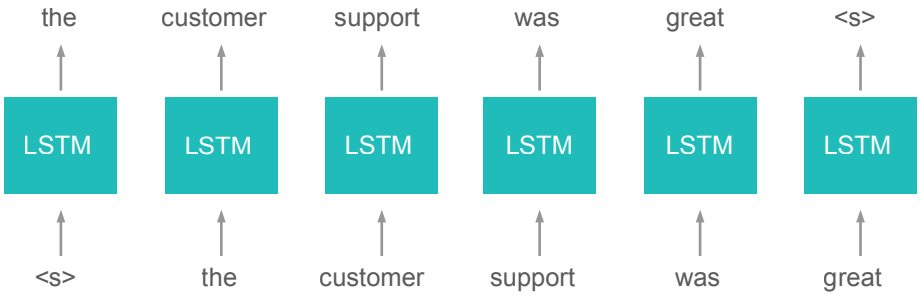

Figure 4: The LSTM architecture can understand long strings of words contextually, leading to better sentiment outputs.

Deep learning also offers the added advantage of inference-based learning. Using RNNs, it's possible to make fine-grained opinion analyses capable of detecting subjective expressions in emails, Tweets, messages, or comments, characterizing their intensity and sentiment, and identifying the object of the sentiment. This can be applied now with great accuracy and speed for opinion mining, natural language processing, or summarization.

Using RNNs, it's possible to make fine-grained opinion analyses capable of detecting subjective expressions in emails, Tweets, messages, or comments...

What We're Doing Now

As the accuracy, flexibility, and reliability of sentiment analysis continue to improve, it is leading to important new applications, particularly with respect to real-time analysis. An obvious example is finding out what people think of a particular product, service, event, or news items using Twitter feeds.

Here at TCS, we're applying artificial intelligence and deep learning to new applications for customers in the financial services industry (real-time sentiment analysis and fraud detection), and continue to explore the enormous potential of deep learning for the digital enterprise.

About The Author

Sunil Karkera

Head of Digital Reimagination™ Studio

Sunil is an engineer by profession, and over the past 20 years, he founded three successful startups in Silicon Valley (wireless products, SaaS solutions, and mobile applications), led information systems groups at News Corporation, and worked and consulted for companies such as Nest, Westinghouse, Dell-Sonicwall, Lockheed, and Siebel. He is also a trained typographer.

About Artificial Intelligence at TCS and our Digital Reimagination™ Studio

At TCS, we operate with the belief that the future of business will be driven by five powerful digital forces: mobility and pervasive computing, the cloud, Big Data, social media, and artificial intelligence (AI). We are applying AI, notably deep learning, to all kinds of applications from autonomous vehicles to the analysis of sensor data from the Internet of Things, from fraud detection to natural language processing and conversational agents.

The TCS Digital Reimagination™ Studio is dedicated to helping businesses create fundamentally new experiences by reimagining industries through creative thinking. The Studio brings a start-up culture to large enterprise clients by leveraging the best of world-class creative, design, engineering, and business domain experts. The result is business transformation through rapid product prototyping and extremely agile collaboration.

Contact

To learn more, contact the TCS Digital Reimagination™ Studio at digital.enterprise@tcs.com

Subscribe to TCS White Papers

TCS.com RSS: http://www.tcs.com/rss_feeds/Pages/feed.aspx?f=w

Feedburner: http://feeds2.feedburner.com/tcswhitepapers

About Tata Consultancy Services Ltd (TCS)

Tata Consultancy Services is an IT services, consulting, and business solutions organization that delivers real results to global business, ensuring a level of certainty no other firm can match. TCS offers a consulting-led, integrated portfolio of IT and IT-enabled infrastructure, engineering, and assurance services. This is delivered through its unique Global Network Delivery Model™, recognized as the benchmark of excellence in software development. A part of the Tata Group, India's largest industrial conglomerate, TCS has a global footprint and is listed on the National Stock Exchange and Bombay Stock Exchange in India.

For more information, visit us at www.tcs.com

TCS Design Services I M I 11 I 16