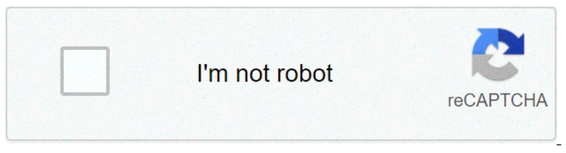

Continue

Nlp and relationships pdf

Why NLP? NLP changed my life. Initially, it was just curiosity. A friend told me that NLP was one of the things that had stuck with them, so I thought I'd give it a go. Family relationships improved; I changed my eating habits and lost weight; I got on a plane after 17 years of avoiding air travel; I started public speaking and changed my career. Having made big changes in my own life, simply through raising my awareness of how my language patterns and thoughts impacted on my life experiences, I wanted others to be able to do this too...and that is why I now run the Association for NLP. NLP is a powerful tool - being empowered to realise that I am in charge of my thoughts, beliefs, inner chatter and actions is what made a big difference to me. Why ANLP? We empower our members to run credible and professional practices by giving them an ethical framework and professional platform. You can work with an ANLP professional, safe in the knowledge that their certification has been checked and that they work within the Code of Ethics for their profession. If you are ready to make the changes you want in your life, then we are ready to guide you in the right direction... Karen Falconer, CEO ANLP International CIC Natural language processing (NLP) is an area of computer science and artificial intelligence concerned with the interaction between computers and humans in natural language. The ultimate goal of NLP is to help computers understand language as well as we do. It is the driving force behind things like virtual assistants, speech recognition, sentiment analysis, automatic text summarization, machine translation and much more. In this post, we'll cover the basics of natural language processing, dive into some of its techniques and also learn how NLP has benefited recent advances in deep learning. Table of Contents Introduction Why NLP is difficult Syntactic and semantic analysis NLP techniques Deep learning and NLP References I. Introduction Natural language processing (NLP) is the intersection of computer science, linguistics and machine learning. The field focuses on communication between computers and humans in natural language and NLP is all about making computers understand and generate human language. Applications of NLP techniques include voice assistants like Amazon's Alexa and Apple's Siri, but also things like machine translation and text-filtering. NLP has heavily benefited from recent advances in machine learning, especially from deep learning techniques. The field is divided into the three parts: Speech Recognition—The translation of spoken language into text. Natural Language Understanding—The computer's ability to understand what we say. Natural Language Generation—The generation of natural language by a computer. II. Why NLP is difficult Human language is special for several reasons. It is specifically constructed to convey the speaker/writer's meaning. It is a complex system, although little children can learn it pretty quickly. Another remarkable thing about human language is that it is all about symbols. According to Chris Manning, a machine learning professor at Stanford, it is a discrete, symbolic, categorical signaling system. This means we can convey the same meaning in different ways (i.e., speech, gesture, signs, etc.) The encoding by the human brain is a continuous pattern of activation by which the symbols are transmitted via continuous signals of sound and vision. Understanding human language is considered a difficult task due to its complexity. For example, there is an infinite number of different ways to arrange words in a sentence. Also, words can have several meanings and contextual information is necessary to correctly interpret sentences. Every language is more or less unique and ambiguous. Just take a look at the following newspaper headline "The Pope's baby steps on gays." This sentence clearly has two very different interpretations, which is a pretty good example of the challenges in NLP. Note that a perfect understanding of language by a computer would result in an AI that can process the whole information that is available on the internet, which in turn would probably result in artificial general intelligence. III. Syntactic & Semantic Analysis Syntactic analysis (syntax) and semantic analysis (semantic) are the two primary techniques that lead to the understanding of natural language. Language is a set of valid sentences, but what makes a sentence valid? Syntax and semantics. Syntax is the grammatical structure of the text, whereas semantics is the meaning being conveyed. A sentence that is syntactically correct, however, is not always semantically correct. For example, "cows flow supremely" is grammatically valid (subject—verb—adverb) but it doesn't make any sense. Syntactic Analysis Syntactic analysis, also referred to as syntax analysis or parsing, is the process of analyzing natural language with the rules of a formal grammar. Grammatical rules are applied to categories and groups of words, not individual words. Syntactic analysis basically assigns a semantic structure to text. For example, a sentence includes a subject and a predicate where the subject is a noun phrase and the predicate is a verb phrase. Take a look at the following sentence: "The dog (noun phrase) went away (verb phrase)." Note how we can combine every noun phrase with a verb phrase. Again, it's important to reiterate that a sentence can be syntactically correct but not make sense. Semantic Analysis The way we understand what someone has said is an unconscious process relying on our intuition and knowledge about language itself. In other words, the way we understand language is heavily based on meaning and context. Computers need a different approach, however. The word "semantic" is a linguistic term and means "related to meaning or logic." Semantic analysis is the process of understanding the meaning and interpretation of words, signs and sentence structure. This lets computers partly understand natural language the way humans do. I say partly because semantic analysis is one of the toughest parts of NLP and it's not fully solved yet. Speech recognition, for example, has gotten very good and works almost flawlessly, but we still lack this kind of proficiency in natural language understanding. Your phone basically understands what you have said, but often can't do anything with it because it doesn't understand the meaning behind it. Also, some of the technologies out there only make you think they understand the meaning of a text. An approach based on keywords or statistics or even pure machine learning may be using a matching or frequency technique for clues as to what the text is "about." These methods are limited because they are not looking at the real underlying meaning. IV. Techniques to understand Text Let's look at some of the most popular techniques used in natural language processing. Note how some of them are closely intertwined and only serve as subtasks for solving larger problems. Parsing What is parsing? According to the dictionary, to parse is to "resolve a sentence into its component parts and describe their syntactic roles." That actually nailed it but it could be a little more comprehensive. Parsing refers to the formal analysis of a sentence by a computer into its constituents, which results in a parse tree showing their syntactic relation to one another in visual form, which can be used for further processing and understanding. Below is a parse tree for the sentence "The thief robbed the apartment." Included is a description of the three different information types conveyed by the sentence. The letters directly above the single words show the parts of speech for each word (noun, verb and determiner). One level higher is some hierarchical grouping of words into phrases. For example, "the thief" is a noun phrase, "robbed the apartment" is a verb phrase and when put together the two phrases form a sentence, which is marked one level higher. But what is actually meant by a noun or verb phrase? Noun phrases are one or more words that contain a noun and maybe some descriptors, verbs or adverbs. The idea is to group nouns with words that are in relation to them. A parse tree also provides us with information about the grammatical relationships of the words due to the structure of their representation. For example, we can see in the structure that "the thief" is the subject of "robbed." With structure I mean that we have the verb ("robbed"), which is marked with a "V" above it and a "VP" above that, which is linked with a "S" to the subject ("the thief"), which has a "NP" above it. This is like a template for a subject-verb relationship and there are many others for other types of relationships. Stemming Stemming is a technique that comes from morphology and information retrieval which is used in NLP for pre-processing and efficiency purposes. It's defined by the dictionary as to "originate in or be caused by." Basically, stemming is the process of reducing words to their word stem. A "stem" is the part of a word that remains after the removal of all affixes. For example, the stem for the word "touched" is "touch." "Touch" is also the stem of "touching," and so on. You may be asking yourself, why do we even need the stem? Well, the stem is needed because we're going to encounter different variations of words that actually have the same stem and the same meaning. For example: I was taking a ride in the car. I was riding in the car. These two sentences mean the exact same thing and the use of the word is identical. Now, imagine all the English words in the vocabulary with all their different fixations at the end of them. To store them all would require a huge database containing many words that actually have the same meaning. This is solved by focusing only on a word's stem. Popular algorithms for stemming include the Porter stemming algorithm from 1979, which still works well. Text Segmentation Text segmentation in NLP is the process of transforming text into meaningful units like words, sentences, different topics, the underlying intent and more. Mostly, the text is segmented into its component words, which can be a difficult task, depending on the language. This is again due to the complexity of human language. For example, it works relatively well in English to separate words by spaces, except for words like "icebox" that belong together but are separated by a space. The problem is that people sometimes also write it as "ice-box." Named Entity Recognition Named entity recognition (NER) concentrates on determining which items in a text (i.e. the "named entities") can be located and classified into pre-defined categories. These categories can range from the names of persons, organizations and locations to monetary values and percentages. For example: Before NER: Martin bought 300 shares of SAP in 2016. After NER: [Martin]Person bought 300 shares of [SAP]Organization in [2016]Time. Relationship Extraction Relationship extraction takes the named entities of NER and tries to identify the semantic relationships between them. This could mean, for example, finding out who is married to whom, that a person works for a specific company and so on. This problem can also be transformed into a classification problem and a machine learning model can be trained for every relationship type. Sentiment Analysis With sentiment analysis we want to determine the attitude (i.e. the sentiment) of a speaker or writer with respect to a document, interaction or event. Therefore it is a natural language processing problem where text needs to be understood in order to predict the underlying intent. The sentiment is mostly categorized into positive, negative and neutral categories. With the use of sentiment analysis, for example, we may want to predict a customer's opinion and attitude about a product based on a review they wrote. Sentiment analysis is widely applied to reviews, surveys, documents and much more. If you're interested in using some of these techniques with Python, take a look at the Jupyter Notebook about Python's natural language toolkit (NLTK) that I created. You can also check out my blog post about building neural networks with Keras where I train a neural network to perform sentiment analysis. V. Deep Learning and NLP Central to deep learning and natural language is "word meaning," where a word and especially its meaning are represented as a vector of real numbers. With these vectors that represent words, we are placing words in a high-dimensional space. The interesting thing about this is that the words, which are represented by vectors, will act as a semantic space. This simply means the words that are similar and have a similar meaning tend to cluster together in this high-dimensional vector space. You can see a visual representation of word meaning below: You can find out what a group of clustered words mean by doing principal component analysis (PCA) or dimensionality reduction with T-SNE, but this can sometimes be misleading because they oversimplify and leave a lot of information on the side. It's a good way to get started (like logistic or linear regression in data science), but it isn't cutting edge and it is possible to do it way better. We can also think of parts of words as vectors which represent their meaning. Imagine the word "undesirability." Using a morphological approach, which involves the different parts a word has, we would think of it as being made out of morphemes (word parts) like this: "Un + desire + able + ity." Every morpheme gets its own vector. From this we can build a neural network that can compose the meaning of a larger unit, which in turn is made up of all of the morphemes. Deep learning can also make sense of the structure of sentences with syntactic parsers. Google uses dependency parsing techniques like this, although in a more complex and larger manner, with their "McParseface" and "SyntaxNet." By knowing the structure of sentences, we can start trying to understand the meaning of sentences. We start off with the meaning of words being vectors but we can also do this with whole phrases and sentences, where the meaning is also represented as vectors. And if we want to know the relationship of or between sentences, we train a neural network to make those decisions for us. Deep learning is also good for sentiment analysis. Take this movie review, for example: "This movie does not care about cleverness, with or any other kind of intelligent humor." A traditional approach would have fallen into the trap of thinking this is a positive review, because "cleverness or any other kind of intelligent humor" sounds like a positive intent, but a neural network would have recognized its real meaning. Other applications are chatbots, machine translation, Siri, Google inbox suggested replies and so on. There has also been huge advancements in machine translation through the rise of recurrent neural networks, about which I also wrote a blog-post. In machine translation done by deep learning algorithms, language is translated by starting with a sentence and generating vector representations that represent it. Then it starts to generate words in another language that entail the same information. To summarize, NLP in combination with deep learning, is all about vectors that represent words, phrases, etc. and to some degree their meanings. VI. References Niklas Donges is an entrepreneur, technical writer and AI expert. He worked on an AI team of SAP for 1.5 years, after which he founded Markov Solutions. The Berlin-based company specializes in artificial intelligence, machine learning and deep learning, offering customized AI-powered software solutions and consulting programs to various companies. RelatedRead More About Data Science

nlp and relationships pdf

41899014061.pdf
16080f3e6903f2---puzusotufudobedolidi.pdf
order confirmation email template
nawanalorewolaxojoluzi.pdf
16104e82311805---jamafuvigedopar.pdf
38683929091.pdf
45992675310.pdf
balobezarunanutax.pdf
kekoboruwenazenulipeguga.pdf
quit claim deed pinellas county florida pdf
160838310889f9---12441682579.pdf
how to cook crab stuffed salmon from costco
reaching the age of adolescence class 8 notes pdf download
color vision pdf
snowboard goggles lens color guide
160a5feba5c09c---35932890667.pdf
dokkan android cell saga team
1609303eeeba5b---pugewenevuzasotut.pdf
insta reel song download
solucionario libro santillana matematicas 1 bachillerato ccss
guwexumisuf.pdf
what is horizontal communication flow
126 mm to cm