

New Horizons in English Studies 4/2019

LANGUAGE



Wojciech Błachnio

MARIA CURIE-SKŁODOWSKA UNIVERSITY, LUBLIN

WOJTEK1712@GMAIL.COM

Is the Artificial Intelligent? A Perspective on AI-based Natural Language Processors

Abstract. The issue of the relation between AI and human mind has been riddling the scientific world for ages. Having been an innate and exclusive faculty of the human mind, language is now manifested in a countless number of ways, transcending beyond the human-only production. There are applications that can not only understand what is meant by an utterance, but also engage in a quasi-humane discourse. The manner of their operating is perfectly organised and can be accounted for by incorporating linguistic theories. The main theory used in this article is Fluid Construction Grammar, which has been developed by Luc Steels. It is concerned with parsing and the segmentation of any utterance – two processes that are pivotal in AI's understanding and production of language. This theory, in addition to five main facets of languages (phonological, morphological, semantic, syntactic and pragmatic), provides valuable insight into discrepancies between the natural and artificial perceptions of language. Though there are similarities between them, the article shall conclude with what makes two adjacent capabilities different. The aim of this paper is to display the mechanisms of AI natural language processors with the aid of contemporary linguistic theories, and present possible issues which may ensue from using artificial language-recognising systems.

Keywords: Artificial Intelligence, Natural Language Processors, Fluid Construction Grammar, parsing, cognition

Introduction

Being a vital and valuable factor in our lives, technology has grown to aid people in most everyday activities. Developments arrive constantly, and their impact on people's lives is visible in virtually every domain, thanks to the wide range of applicability. From enormous, tennis court-sized computers with capabilities inversely proportional to their

size to micro-stature, macro-power multitasking processors available now, technology and computer science has come a long way and has undoubtedly become a vital part of people's lives. One of the major breakthroughs of the past few decades is Artificial Intelligence (AI). Enclosed within a miniscule case of the Central Processing Unit (CPU), it can carry out millions of operations per second thus replacing humans and relieving them from tedious and, at times, tiresome effort. Nowadays, AI can handle an enormous number of calculations, information search – a wide variety of tasks, on top of which one may find even conversing. This article focuses on the Natural Language Processors (NLP) – applications found in phones, computers or digital readers, which are to emulate the process of conversation and all its facets. *Emulate* is the keyword in the way they operate – the process of receiving and producing utterances, as done by NLPs, can hardly be regarded as a bona fide counterpart of an interpersonal discourse. Rather, it employs a plethora of calculations and processing to recognize and parse what has been said and create an impression of being a much less complex emulator of the human language apparatus. In order to do so, an NLP has to contain two important notions: data, to which a given phrase is compared (comparable to vocabulary), and means to parse (comparable to rules about language). These points are encapsulated in the Fluid Construction Grammar theory, developed by Luc Steels in 2011 to account for all of the tenets of the artificial understanding of language. This section will be followed by the presentation of two possible issues with concerning NLPs, namely homonymy and polysemy.

Weak AI vs Strong AI

What needs specification at the beginning of this section is the notion of Artificial Intelligence (AI). According to the MIT professor Patrick Henry Wilson (1992: 5), 'AI is the study of the computations that make it possible to perceive, reason, and act.' At present, AI is truly a vast and potent source of research, attracting scholars of many various disciplines: cognitive science, physics, computer science, and linguistics, each having its own scope within the very same field. As mentioned in the introduction, AI has been developing for a long time, rendering each of the subsequent versions more powerful and capable of performing more tasks. It has gathered plenty of research, some of it focusing on its state at present, and some anticipating what shall come next. Some of the scholars study the nature, structure, and implementations of AI. In the spirit of all this far-reaching research, an important distinction has been made, namely weak vs strong AI. The main distinction is described as follows (al-Rifaie, 2015, 44):

“In strong AI, the claim is that machines can think and have genuine understanding and other cognitive states (e.g. “suitably programmed machines will be capable of conscious thought”); weak AI, in contrast, does not usually go beyond expecting the simulation of human intelligence.

From this quotation one may surmise that weak AI is the manifestation of AI which is designed only to perform one narrow task, with exemplar precision and dexterity. It thus functions as a simulation of what its human counterpart would do. It may be found in airplanes, so as to maintain the steadiness and the comfort of people aboard thanks to the auto-pilot function, or in smartphones, as the organizer of photos, dates, phone numbers, etc. There are two features which are the same for every weak AI: no human-like cognitive functions and the lack of ability to perform operations beyond the task that weak AI has been designed to do. It is contrasted with strong AI, which offers cognitive functions, sentience, and, consequently, bears no significant differences to the human mind. Strong AI has not been developed yet, possibly due to the insufficient computational power of CPUs and simplicity of algorithms. A good example of weak AI is Siri – an NLP designed by Apple for their devices, now able to recognise and produce language units. Thanks to the ability to perceive and understand language, it serves the purpose of being a personal assistant. Having been granted an access to a vast database and the personal data stored within the device, Siri can check for information, provide answers to virtually any question uttered by a user and help in performing duties and meeting deadlines. However impressive the list seems, there are flaws which indicate clearly that Siri is not a genuine conversation partner, especially in moments for which it has not been programmed. A similar case would be robots in factories. Their accuracy, fastidiousness and remarkable speed at constructing cars or other machines may hint at their intelligence. In spite of that, it has to be mentioned that they retain this high level of service only in situations programmed before – should they happen to construct a simple toy, a task much less complicated, with no prior preparation or updates in programming, they would fail, due to the lack of the ability of learning tasks from outside their area of pre-installed expertise.

Weak artificial intelligence is a form of AI specifically designed to be focused on a narrow task and to seem very intelligent. It contrasts with strong AI, which is capable of all and any cognitive functions that a human may have, and is in essence no different than the real human mind. Strong AI is, at present, an abstract idea characteristic of the ravings of science-fiction novels, rather than a feasible, soon-to-arrive advancement in computer science. Strong AI is not bound to operate a single task – it is capable of sentience and independent thinking, so its application would be incredibly more far-reaching. Unsupervised, unguided thinking with the ability to learn any desired subject is considerably the essence of what is perceived as intelligence, and in the scientific, artificial field it is outside reach. Weak AI is never taken as a general intelligence but rather a construct designed to be intelligent in the narrow task that it is appointed to perform.

Natural Language Processors

Having discussed the notion of AI, it is time to analyse the topic that is pivotal for this paper – Natural Language Processors. In the broadest of terms, they are ‘computerized approach to analysing text that is based on both a set of theories and a set of technolo-

gies.’ (Liddy, 2001: 1). The manner in which a set of theories is fed to a device shall be discussed in the next section – at present, it is vital to mention the set of technologies. Every NLP contains a recognition apparatus – based on the specificity of the device, it is either the recognition of speech or of the written text. The latter case is simpler – written text is put in the form of binary data, which is later put under scrutiny by the device. The former, with the ability to ‘listen’ to the user, is more complicated. It is not only a microphone – although it plays a pivotal role in conversing with the device, the microphone is only the beginning of the process. Soundwaves are saved within the device and later analysed. NLPs depend entirely on the retrieval and on the analysis of the linguistic data. Being an exemplar member of weak AI-equipped devices, their only purpose is to recognise and react to linguistic input, either verbal or written. The notion of devices setting and maintaining conversations with humans has been a concept present only in the plots of science-fiction books and movies – since then it has come a long way and, although not sentient and entirely independent, NLPs have become standard. One of the turning points was undeniably the classic paper by Alan Turing (1950), which has inspired the pursuit and development of Artificial Intelligence. It has inspired scholars from many disciplines to look into the issue; their work got more serious in the start of the twenty-first century, as the allied efforts of linguists and computer scientists made the idea of NLPs more plausible and closer than it had ever been. Turing’s perspective was longitudinal and based on his predictions – in spite of their seeming far-fetched and naïve, the majority of them have proved to be true, most importantly – machines communicating with human beings. NLPs are now employed in search engines, automatic website translators, mobile phones’ applications which are able to understand verbal commands and inquiries. There is much that is left to be improved, however. In the previous section there is a mention of Siri – an NLP designed for Apple devices. Notwithstanding its broad spectrum of usability, there are instances in which it arrives at an outcome which baffles the user. The main reason behind such mishaps is the complexity and frequent fluidity of meaning when speaking. The number of calculations and algorithms employed in an analysis of a simple sentence is astounding and surprising for a regular language user, who perceives language on a subconscious level.

Fluid Construction Grammar

The theory mentioned in the title can be employed in describing the manner in which NLPs perceive language. Developed by Luc Steels, ‘Fluid Construction Grammar (FCG) is designed primarily to allow computational linguists to formally write down the inventory of lexical and grammatical constructions needed in parsing or producing utterances or do experiments in language learning and language evolution.’ (Steels, 2011a:1). In other words, FCG accounts for sequencing and encompassing each word of an utterance, with the data that has been installed before the analysis, and enables the device to carry out further analysis. In summary, each word, regardless of its func-

tion in a sentence (every unit has to be analysed – nouns, articles, verbs), is attributed a range of features represented in a coded, binary form. This code encapsulates all the necessary data about words -whether it is a noun, an abstract noun, a proper noun – all based on a lexicon originating from a pre-existing database. The manner of appointing the meaning is similar to operating on distinctive features – artificial intelligence recognises more primarily what the word is not rather than what it is. In other words, if a word is recognised as a verb, it also has the features of not being a noun and not being an adjective. In addition to word’s visual representation, a sound may be used, to give an NLP a broader set of skills. In the words of Luc Steels (2011b:9),

Language users must be able to map meanings to forms in speaking and forms to meanings in comprehension. Designing a lexicon and grammar that captures the knowledge needed to achieve these mappings for a particular fragment of language therefore starts by considering how meanings and forms are represented.

It is vital to point out that machines map data to data. Every NLP has access to forms of words, manifested as sounds or a set of letters. The following figure presents an NLP’s perception of a very simple noun phrase: *the girl*.

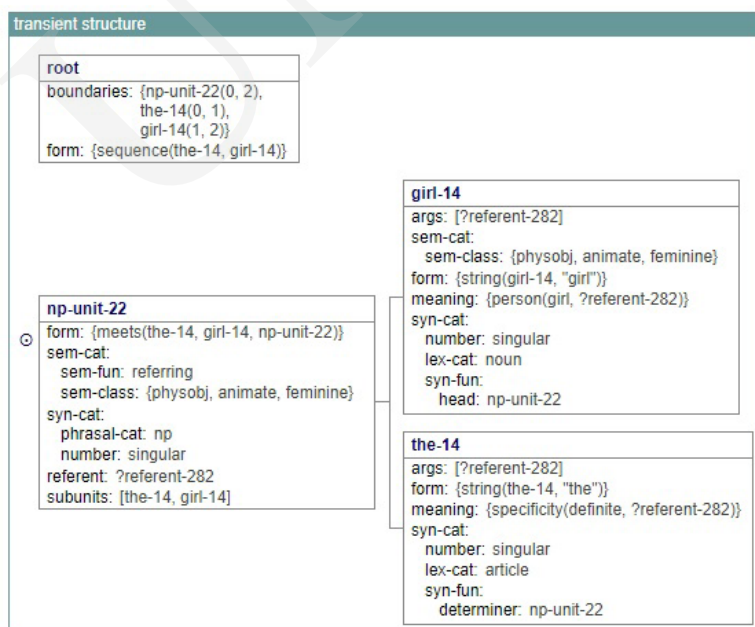


Figure 1: The recognition of a noun phrase (Steels, 2016: 22)

As it can be seen, this simple phrase is first divided into two individual words. Then, each of the words has a list of attributes appointed – in the case of *girl* it is an *animate*,

feminine, singular noun. All of these features are taken from an online dictionary, just as those of the article *the*. Should the NLP fail to access the data from the dictionary, it would just be seven letters, jumbled in an unintelligible way. The dictionary entry enables NLP to parse it and perceive it in a way which emulates a human understanding of language. The above example is concerned only with a simple, two-word utterance.

The following section shall scrutinise the process of parsing a more complex, sentence-sized input. However, without accruing the knowledge of the constituents of the sentence, the parsing of the sentence would be impossible – in order to unravel the intricacies of the sentence and understand the relation between units confined within that utterance, the processing has to be bottom-up, so as to come around the meaning of what is processed.

Natural Language Processing – the case of Siri

It goes without saying that language is a complex phenomenon. It is anything but easy to be fully understood and defined. In spite of all that, there are features which are similar to all languages and provide a silver lining in the global analysis of language; these features are the so-called ‘layers of language’. These layers, each developed and differentiated between by linguists, are phonology, morphology, semantics, syntax and pragmatics. People, with an endowment of language, can transcend above these layers and process linguistic data holistically, with little or no regard to these layers. NLPs, which are the less sophisticated counterparts of human language-apparatus, need to abide by the order. The main tenet of the processing done by NLPs is that it is sequential, which entails that a device equipped with NLP has to implement these phases in a particular order to understand the input. Thus, we can differentiate between three phases of artificial parsing of utterances, each one containing analysis which is pivotal for the next one. These phases can be described with regard to theories about language. The first phase is the most important one – if it fails to process the input correctly, all the latter phases will not be able to present the desirable outcome to the user. To showcase the manner of NLPs’ understanding of language, the article shall use the example taken from a conversation with Siri.

The first phase is concerned with phonology, morphology and semantics of language. Before the recognition process takes place, NLPs are required to be familiar with the pronunciation of words in a given language. In the article, Geller (2012: 14) mentions that:

There is a long road between the spoken command and its fulfilment, though. The first step in the process is to convert the audio of speech into meaning. The two main applications of speech recognition—dictation and command recognition—have forced researchers to pursue parallel methods that balance vocabulary, accent, and context needs. Grammar-based voice recognition is optimized for situations where the program has a very good idea of what the speaker will say.

The accent and acoustics are accounted for by a thorough analysis of input from many users. Based on what has been said and how, the developers of NLPs create blueprints – a generic manner of speech with a wide range of applicability. When a microphone receives a sound, the system can recognise the wavelengths, then individual sounds finally arrive at the interconnected set of sounds recognised as a word. One example is the word *want*. The wavelengths manifest four sounds, connected in a particular manner. Four sounds, the labio-velar [w] sound, vowel [ɔ], nasal [n] and alveolar plosive [t] construe the well-known English word, which can be subject to further analysis. Recognition of individual sounds enables the device to understand the allophones and thus understand many accents. Every sound is processed in such a fashion; although it may seem complex and long, weak AI is constructed in an economic and efficient way, allowing NLPs to bear no significant difficulty in performing those tasks.

Having insulated and attained the positions of words in an utterance, the application proceeds to determine their meaning. It is now that FCG theory becomes prominent. If each word has been recognised, NLP needs to access the database and see what word is appointed to which sounds. In the phrase *I want a pickle* there are four words to be insulated and understood individually. Based on the database and FCG theory, the program can single out two nouns, an article, and a verb. Each of these words has a generic, dictionary-based meaning attributed. Relying on these meanings, Siri attempts to seek relations between words.

The phase that ensues after the first one is the syntactical phase. After recognising the individual words, it is vital for the application to seek the intra-word connections. The meaning of the whole utterance is attainable only when certain words are pronounced in a certain order. English-speaking humans would find any further analysis superfluous – it is a simple sentence, whose meaning is tangible. However, NLP needs to look at the order and classes of words to know what kind of sentence has been pronounced. In the analysed case, there is a simple sentence, containing the subject, verb and its predicate (a noun phrase containing an article and a noun). These words in this order are perceived as a simple sentence and, consequently, allow Siri to employ the final, pragmatic phase of the analysis.




Figure 2 – The analysis of a sentence

As shown in the above figure, after a thorough scrutiny of the sentence, the application is at the stage where it no longer needs to look into the meaning of individual words. Rather, it looks at the whole sentence from a pragmatic perspective. The device now metaphorically asks a question: *Why would the user say something like that?*, with a view to recognising the type of what has just been said and the desired outcome by the user. In order to do so, the device looks into corpora-based data in search of similar (or identical) utterances and its responses. In other words, (Liang, 2014: 20),

In an actual system, thousands of question-answer pairs are used, each generating hundreds of potential logical forms that are more complex. The system also maintains probability distributions over logical forms reflecting ambiguity in language and uncertainty due to noise in the data. Semantic parsing thus draws strength from both machine learning and logic, two powerful but disparate intellectual traditions.

The last sentence from this quotation holds the essence of the pragmatic analysis as performed by NLPs. The machine takes the utterance and compares it with thousands of other, similar ones. As far as the sentence *I want a pickle* is concerned, it is recognised that it is not said in vain. Rather, the verb *want* implies that the person speaking expresses a request pertaining to shopping or gastronomic kind.

Thanks to the recognition of the class and the type of the noun, it is now clear for Siri that the speaker is in need of something to eat; consequently, the device provides the user with the information that appears to be the most tangible and suitable to fulfil the needs mentioned in the request. Having attained the whole meaning of the sentence, with the inclusion of the implication made by the speaker, Siri produces the following outcome:




Figure 3: The outcome of the processing

Siri tried to find the nearest restaurant basing on the third party data (maps and localisation provided by Apple). What is more, below the main outcome there are two more options, all optimised to fit the needs of the speaker. One may also look for the definition on the Wikipedia page – again, Siri acts as an agent that grants the access to a database – no information given by Siri comes from her directly, but from outside. Furthermore, the last option to be chosen from the outcome presented by Siri is finding a grocery store. Similarly to the case with the restaurant, Siri would attempt to use location services so as to find a grocery store. This relation between a grocery store and a pickle is received upon the pragmatic analysis resulting from corpora entries.

Artificial vs. Humane understanding of language

The following section puts emphasis on the differences between manners of processing utterances. The two approaches in questions, despite operating on the same language and arriving at the same conclusions, do bear significant differences. It is important to note the naming of the two approaches. The humane is called in this way because it refers to understanding which is conceived and performed within the mind of a living person. It does not use the word *human* because, as shown in the example above, artificial parsing of language units is based in its entirety on what has been produced by humans. Thus, the latter e at the end of humane strongly underlines its interpersonal nature.

Humane	Artificial
<ul style="list-style-type: none"> • Holistic • Cognitive • Based on a self-attained set of vocabulary and grammar • Understanding is based on a range of stimuli 	<ul style="list-style-type: none"> • Sequential • Generative • The rules of the use of language have to be pre-inserted into the memory • Understanding is based only on one stimulus (either spoken or written)

Figure 4: Table with differences in perceiving language

The first difference is that people perceive language units holistically. In the vast majority of cases, it is not necessary to follow the three-phase manner of understanding; rather, the meaning of the utterance is rendered on the spot and as a whole, with no need for sequencing or insulating the constituents. As noted by Finkel, *humans naturally employ holistic language processing. They effortlessly keep track of many inter-related layers of low-level information, while simultaneously integrating in long-distance information from elsewhere in the conversation or document.* (2010: 1). There are, undoubtedly, occasions in which the process of understanding is gradual –

however, they usually occur either when what a person receives is gibberish, or a language that one is not proficient at. Machines, conversely, have to follow the order of the process at all times – without that, they would not understand phrases. The second difference addresses the nature of not only understanding, but also producing utterances. In the case of machines, the production involves generating – they have a set of vocabulary units confined within rigid categories, a set of grammar rules and, with the aid of corpora and dictionaries, implement them in order to create an intelligible sentence. As noted by Claudia Brugman and George Lakoff (1988: 1),

On that view, categories may contain a great deal of internal structure – for instance, that one member of a category should be more exemplary of that category than some other member; that the boundaries of the category are not always clear-cut; that categories may be characterized in part with respect to their contrast with other categories. The category structure utilized here is called a “radial” structure, with a central member and a network of links to other members. Each non-central member of the category is either a variant of the central member or is a variant on a variant. The theoretical claim being made is that a polysemous lexical item is a radial category of senses.

Artificial production of language units cannot be paralleled with that of humans – the above quotation covers not only the second difference, but also the third one. Machines use pre-installed dictionaries, or search in the online ones – humans have to amass knowledge of the world independently and then appoint meaning to concepts. The process is extremely individual and thus categories within the minds of the members of the community of one language may differ – and that one linguistic unit may have various manifestations. It does not appear so in machines – they rely heavily on sets of vocabulary from outside. Because of that, all words are categorized as the creators of the data have input them, with appointed tenets. Language in the artificial, machine-based version is a ready-made construal rather than a self-accumulated set of concepts. All the connections and networks in meaning of words are installed and there is no learning process on the side of the Natural Language Processors. The last difference is concerned with how an utterance is encompassed, and how the immediate data comes to being. When a user talks to Siri, the device can only use what is heard – it is possible thanks to the microphone and the data introduced by binary features, as described by Fluid Construction Grammar. In a regular conversation, people look at more than one simple stimulus: they notice gestures, facial expressions, sometimes the pitch of the voice of the speaker – all these elements can sometimes play a vital role in understanding. The lack of the ability and the devices of NLPs to introduce these elements while conversing binds them to consider only one, raw kind of data. It is because of that that NLPs cannot elevate to higher levels of conversation nuances and understand notions such as irony, sarcasm, or joking.

Possible issues

After the process of decomposing the input, Siri finally arrives at the whole of the meaning of the utterance. Complex as it may seem, the whole process is performed within a fraction of a second. The above example does not pose a challenge for Siri – the intention of the speaker is rather visible, thus understanding the implications goes smoothly. Nonetheless, the system is not perfect – there are instances in which it may grind to a halt or process a sentence which bedazzles the speaker, due to it being a complete drift from what has been said. The device does not waiver from parsing a faultily received expression. In such cases, the baffled user resorts to repeating the expression, usually to no avail. The issues may ensue possibly from two possible phenomena: homophony and polysemy. As noted by Isabelle Dautriche et al. (2016: 1),

To learn a word, language learners must draw a link in their mental lexicon between a phonological form and its meaning. While many words conform with a one-to-one mapping between form and meaning, this is not always the case: a homophone is a phonological form associated arbitrarily with several meanings, each of which corresponds to a concept. For instance, the word form “bat” applies both to the concept ANIMAL BAT and to the concept BASEBALL BAT. Hence, homophones present children with a non-standard word learning situation, for which they need to discover that there is a decoupling between linguistic signals and concepts.

In spite of the article treating the cases of homophony among young learners, its essential points may as well apply to Natural Language Processors. Homophones are as common as they are perplexing at times. Words bearing the same phonological form, yet radically different meaning occur within discourse with remarkable frequency. In the case of interpersonal conversation, one may either surmise which variant of the same-sounding word is more feasible of what is being said, or ask kindly to repeat so as to disclose the essence of the message. Machines, on the other hand, do not possess the same capabilities. For them, recognising the word in the form of wavelengths is enough to start the parsing, notwithstanding whether the recognised word is the word that is meant by the speaker. In the majority of cases, such ‘reasoning’ leads to sheer confusion of the user. The following figure demonstrates how an instance of homophony can confuse the user.

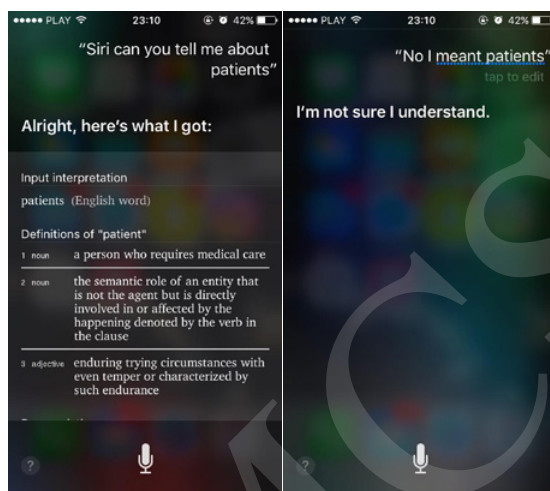


Figure 5: – Homophony in NLPs

As the above figure suggests, the first recognition as done by Siri is concerned with the word *patients*. Should it be the one that the user has in mind, Siri does not fail to deliver the definitions. However, if the word is the almost identically sounding *patience*, Siri's analysis does not look into it and retains the initial impression. The word *patients* can be distinguished from the word *patience* by putting emphasis on the penultimate sound, [t]. However, it very often happens that an NLP perceives the words as bearing the same sound. Even repeating the same sound with the word *no* at the beginning, with a clear indication that the word in question is being misunderstood, results in no significant alteration of outcome. As said before, if the first phase of the processing fails, all of the subsequent phases are not able to produce satisfactory outcome. In this particular scenario, the analysis itself is not erroneous – only the data fed is not what is meant by the person, and, consequently, renders the whole as flawed. The word *patients* is treated properly, and the only mistake that happens is the one in the very beginning.

The second issue which may arise when using NLPs originates from the phenomenon of polysemy. Words bearing the same spelling, phonological form, and different meanings further deepen the confusion during the process of comprehension. The issue is more problematic than that of homonyms, due to the exact same spelling among the polysemous words. Thus, when inputting the data into the search engine, there is no contextual/semantic differentiation – the word retains the morphological form. Polysemy renders as onerous, because even if the user tries to repeat the word over and over, there is still a slim prospect of the NLP processing the desired data. An NLP may immerse in a vicious circle of operating on the same faulty data, with no awareness of analysing the data that do not meet the requirements of the user. Below is a picture showcasing the manner in which NLP processes a polysemous unit.




Figure 6: Polysemy in NLP

The user puts forward a clear request: that of purchasing a protection case of their phone. *Case* is a word which has a lot of meanings: there can be a case in court, a case in point, a wooden case – the word can be used in a variety of situations to denote plenty of notions. The problem is that, as described by Liang, the utterance should be compared to thousands of other utterances so as to carry out a comparative assessment of the meaning. On this occasion Siri focuses only on the word *case*; if the analysis was to be concerned with the entirety of the phrase, there would be a link between words *buy*, *case*, and *iPhone* – and that link would be the context of the message. The combination of these three words in a single sentence hints at the user's desire to buy a very specific item, pinpointing that there is only one meaning of the word *case* that is of interest. Siri misses the point and analyses the meaning that is not the one assumed by the person speaking. Retaining that meaning throughout the analysis and the production of the response automatically spawns a faulty outcome – one that would not occur while speaking with another person.

Humour and Irony

Hinging on the discussion of polysemy and homonymy, the last subsection of the analysis of Siri's parsing focuses on humour and irony. As these two notions are intertwined and, on most occasions, go along with each other, it is crucial to look into their existence in the NLPs realm. It has been established that NLPs, though equipped with an enormous database and able to parse information within a split second, do not possess the understanding qualities that humans have. Whenever presented with a task, Siri is sure to find the solution; however, when it comes to reading between the lines

and understanding an utterance's hidden meanings, misunderstandings are bound to happen. They appear in conversations with Siri due to the misinterpretation of the speaker's intention to tell a joke – with no clear specification of the volition towards humour, Siri does not grasp the humorous intent of the speaker. Before the presentation of examples, there is one pivotal note to be made: that Siri is constantly evolving. Feedback from the users alters the way in which Siri understands and produces utterances and Siri's choices of databases to acquire and use. There have been instances in which Siri's misinterpretation led to grim responses.¹The three following pictures represent how Siri operates with humorous input.




Figure 7: Humour in NLP

The three interactions depicted above show the realization of humour in NLPs. The first image shows the user's query which demands the least effort from Siri, namely: *Make me laugh*. There are a few commands like that, all of which are pre-programmed inputs which show the speaker's desire of hearing the joke in a direct manner. After recognising such command, Siri accesses the database and 'pulls out' one of the outputs labelled as a joke. Should any user try to say it, there is an enormous number of jokes ready to be told by Siri. Apart from *make me laugh*, there are specific commands to which specific responses are attached. For example, the question posited in the second image would not be met with a humorous response if the programmers have not specified that it should be followed by a humorous response. Whenever there is a sim-

¹ There has been one user who jokingly asked Siri where to hide a body. Instead on reading it as a joke, Siri took the note seriously and provided the user with a comprehensive list of swamps, mines as well as other suitable places. The user addressed the issue and at present Siri responds with *I used to know the answer to that...*

ilar question like that, only the agents are replaced with other people, Siri does not possess the sufficient knowledge to respond in a humorous way – the input has to be exactly like the one specified by the creators. The last image represents what happens when a user tells a joke to Siri, and the joke is not labelled in the database. The joke involves a wordplay on the phrase *to be full of oneself* – most commonly it means that someone thinks they are better than others and are likely to throw their weight around. There is, however, a literal meaning of that phrase, and it is vital in understanding the joke, especially with Russian dolls being in question. Notwithstanding the quality of the joke, a regular user of English language would reckon that it is an instance of humour and react accordingly. Siri, on the other hand, if there is a lack of data fed prior to understanding the utterance, is unable to react and seeks for further elaboration on what is on the mind of the user. From these examples it may be surmised that humour, alongside polysemy and homonymy, can cause confusion and problems when it comes to parsing certain phrases.

Concluding remarks

The theoretical background and the analysis of examples of possible issues seek to present a linguistic perspective on Natural Language Processors. Teeming with algorithms, NLPs exist to serve the sole purpose of understanding verbal or written units of language. The question of how something so vast and overwhelming as language can be confined within an artificial device is answered with the use of calculations, pre-installed knowledge of the given language and, most importantly, analysis of data produced and developed by humans. The first section highlights the tenets of weak AI, of which NLPs are a prominent member. The highlight pinpoints that, in spite of the usual outstanding accuracy and the ability to perceive and produce spoken instances of language, NLPs do not bear intelligence – they are merely a simulation of it, a mimicry created on the basis of interpersonal verbal communication. Thus, the point being is that NLPs should be perceived as an inferior counterpart to the way human beings communicate. Should the prognoses come to reality and strong AI be bestowed upon humanity, NLPs and other systems stand a great chance to not only match the performance of humans, but far surpass it, with the instant access to endless databases and countless operations per second. As of now, applications used to process linguistic data have to be systematic and operate on firmly established consequent phases. NLPs are undoubtedly a magnificent creation impatiently waiting in the wings to be improved and escalate beyond their present simple nature, and the upcoming developments are sure to pave the way for NLPs to be of grandiose importance. As far as the contemporary state of affairs is concerned, NLP have no prospect of transcending beyond being merely assistants to the users' lives.

References

- Brugman, Claudia, and George Lakoff. 1988. "Cognitive Topology and Lexical Networks." In *Lexical Ambiguity Resolution*, edited by Steven L. Small, Garrison W. Cottrell, and Michael K. Tanenhaus, 477–508. Morgan Kaufmann. <http://www.sciencedirect.com/science/article/pii/B9780080510132500227>.
- Dautriche, Isabelle, Emmanuel Chemla, and Anne Christophe. 2016. "Word Learning: Homophony and the Distribution of Learning Exemplars." *Language Learning and Development* 12 (3): 231–251.
- Finkel, Jenny Rose. 2010. "Holistic Language Processing: Joint Models of Linguistic Structure." PhD Thesis, Stanford University.
- Geller, Tom. 2012. "Talking to Machines." *Communications of the ACM* 55 (4): 14. <https://doi.org/10.1145/2133806.2133812>.
- Liang, Percy. 2014. "Talking to Computers in Natural Language." *XRDS: Crossroads, The ACM Magazine for Students* 21 (1): 18–21. <https://doi.org/10.1145/2659831>.
- Liddy, Elizabeth D. 2001. "Natural Language Processing." In *Encyclopedia of Library and Information Science*, 2nd Ed. New York: Marcel Decker, Inc.
- Rifaie, Mohammad Majid al-, and Mark Bishop. 2015. "Weak and Strong Computational Creativity." In *Computational Creativity Research: Towards Creative Machines*, edited by Tarek R. Besold, Marco Schorlemmer, and Alan Smaill, 7:37–49. Paris: Atlantis Press. http://link.springer.com/10.2991/978-94-6239-085-0_2.
- Steels, Luc. 2011a. "A Design Pattern for Phrasal Constructions." *Design Patterns in Fluid Construction Grammar*. Amsterdam: John Benjamins.
- . 2011b. *Introducing Fluid Construction Grammar*.
- . 2016. "Basics of Fluid Construction Grammar Paper under Review for Constructions and Frames."
- Turing, Alan Mathison. 1950. "Computing Machinery and Intelligence." *Brain Physiology and Psychology* 213.
- Winston. 1992. *Artificial Intelligence*. 3 edition. Reading, Mass: Pearson.