NATURAL LANGUAGE PROCESSING

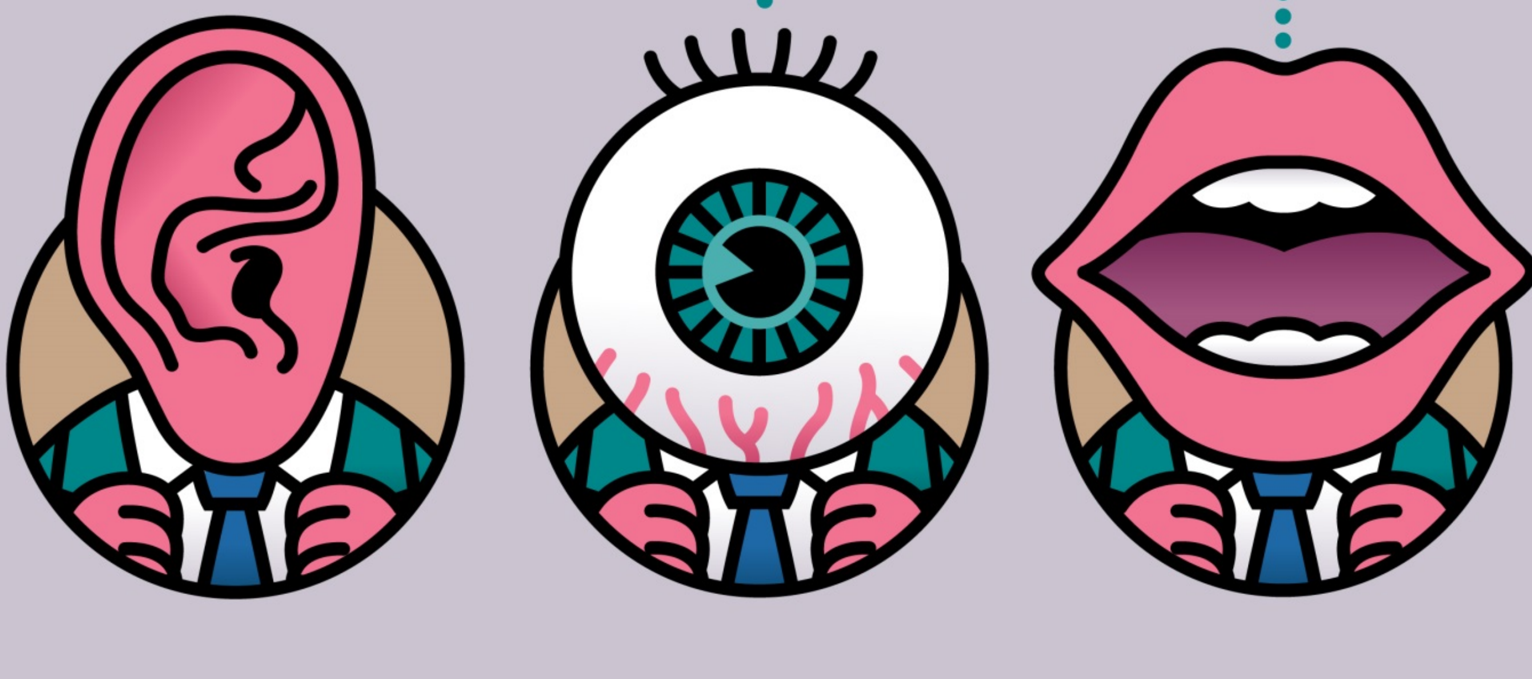

sas

Specialty Answering Service

WWW.SPECIALTYANSWERINGSERVICE.NET

Contents

1. Introduction

Contact centers have approached Natural Language Processing (also known as Natural Language Understanding or Natural Language Technology or just Natural Language) as the ultimate answer to develop high end speech based solutions. NLP aims to improve the customer service experience by permitting the caller to say anything in context rather than constrain him to a set of predefined choices. As the contact center industry trumpets NLT as the future of speech recognition based applications, vendors have started touting the NLP capabilities of their products and research efforts are being focused on solving real world problems using advances in NLP technology.

Natural Language Processing is a branch of Artificial Intelligence that deals with the computerized analysis of naturally occurring text and speech to achieve human-like language processing for a variety of practical applications. NLP makes use of a variety of computational techniques to achieve different types of language analysis. NLP techniques can be applied on any language used by humans and the text input can be either oral or written.

There are multiple types of processing involved when humans communicate using a language. Based on the different types of linguistic analysis capabilities offered by NLP applications, they are referred to as either an NLP system or an NLP-based system. NLP is often a means to achieve a practical goal and the various goals of NLP systems include Information Retrieval, Human-Machine Interaction etc. A full-fledged NLP system will be able to perform the following functions:

1. Paraphrase an input text
2. Translate a given text into another language
3. Provide contextual responses based on the text contents
4. Draw inferences from the text

Today advanced NLP systems are able to perform the first three functions efficiently, however, drawing inferences from a given text is still a challenge in many cases. The goal of an NLP system is dependent on the application for which it is designed. For example, the goal of an NLP based Information Retrieval system shall be to provide accurate information in response to a user's real time requests expressed in an everyday language question.

In contact centers, NLP based systems have been used in speech recognition and IVR applications to provide natural interactions with callers. NLP techniques uses a rule base to define how the system can respond to users and guide them through a transaction rather than dictate to them and restrict them. The defining characteristic of a natural language experience is that it engages callers with a human-like conversation that is intuitive and smooth.

In this paper, we will examine the history and evolution of NLP, the various approaches to NLP and how it can be used to improve the accuracy of speech based applications thereby providing customers a more enriching experience while dealing with your contact center.

2. The Need for Speech Enabled Automation and the Role of NLP

Interactive Voice Response systems are the backbone of contact centers which aim to reduce costs through customer self-service. A well designed IVR system not only helps to curtail operating costs, but also reduces wait time and resolution time for callers by offering smart self service options. This also frees up agent time which can then be spent on resolving more complex customer queries which cannot be handled by self service. In the early days, IVR systems only had touch tone options, thus restricting them to offering simple self service tasks. Today, with speech-enabled IVR systems becoming the norm, complex tasks such as address change are being offered as a self-service transaction, thereby increasing the level of automation in the contact center.

One of the key factors that determine the success of self-service solutions is the accuracy of the speech recognition system during interactions with the customer. Factors such as noisy backgrounds, heavy accents, and responses that are not legible to the system can all make it difficult to complete a transaction.

2. 1. Speaking the Customer's Language

When callers reach out to a contact center with their concerns they typically explain their problem in their own unique vocabulary. Each caller attempts to explain the problem in what they feel is the most simple and clear form. However, for the contact center agents and IVR systems, this means that they have to be able to interpret the customer's need by decoding the vocabulary that is being used. Before NLP was introduced in speech recognition systems, they relied on a predefined grammar (a list of acceptable phrases) that the system is programmed to recognize at each stage of the customer interaction. Grammars ensure that the speech recognition engine interprets "agent", "I want to speak to someone" or "I want an operator" all in the same way and directs the caller to a live agent. This works well as long as the callers stick to phrases defined in the grammar. However, when customers speak something that is out-of-grammar then the speech system is unable to interpret it correctly, leading to frustration on the part of the caller and increased costs on the part of the call center. Studies have shown that out of grammar errors are almost 5 times that of misrecognition errors in most speech application implementations. There are several ways in which natural language processing techniques can be applied to improve the accuracy of the speech recognition system, thereby increasing the automation rate and improving customer satisfaction levels:

2. 2. Enlarge the grammar

The brute-force technique of enlarging the grammar to contain a larger list of entries can mimic natural language capabilities without using actual natural language algorithms. This method is best suited for situations with a limited number of responses. For example, if the IVR system is just looking for a confirmation from the caller, then the grammar can include words such as "yes", "yup", "sure", "yeah" and other variations of a positive response.

However, as the list of possible responses increases, it would become difficult to maintain the system and offer superior performance as the system would have to match what the caller said against the large grammar database. Also, practically, it may not be possible to capture every possible response variant that a caller may use, even if a very large grammar database is maintained. Therefore, in order to offer superior and error free self-service options, natural language capabilities are a must-have.

2. 3. Ability to spot keywords

Since every caller would typically add extraneous phrases to their response, a speech recognition system that can spot the relevant keywords in the speech can improve its accuracy rates. This method is known as robust parsing or keyword spotting. It aims to filter out the filler phrases and interpret just the keywords relevant to the context of the conversation. For example, instead of saying “no” as a response, a caller might say “well, I don’t think so.., no”. It is up to the NLP enabled system to parse this and only use the relevant keyword ‘no’ as the response to the system query. The keyword spotting approach is an effective technique when the dialogs mostly consist of yes/no answer which is what constitutes nearly 70% of speech enabled interactions. By using keyword spotting techniques, callers can be offered more flexibility in how they respond to the IVR interaction, thereby greatly increasing the perceived accuracy of the system. Keyword spotting also reduces retry prompts and repeated confirmations by the system, thus making callers more comfortable during the call. Statistical analysis techniques are used to compare filler phrases collected from transcriptions of previous calls and applying them to analyse the current call. Over a period of time, this can improve recognition accuracy by as much as 30%.

3. Statistical Modeling

High end NLP techniques do not rely on a predefined grammar based matching for conversation recognition. Instead sophisticated statistical modelling techniques are used to identify key phrases and contextual meanings. This allows greater flexibility to the system to ask callers open-ended questions such as “How may I be of assistance?” The system will be able to identify which department the call should be routed to based on an analysis of the user’s responses in a natural language. Thus, when a caller says “I would like to optimize my bill based on my usage pattern’ , the system will be able to correctly route the call to ‘tariff plans’ even though the exact phrase was never used by the caller. However, designing the statistical models is a time-consuming task as the meaning of each transcription needs to be defined manually and previously transcribed phrases need to be tagged properly, Also, if the user gives a completely unexpected response such as “hamburger” to a telephone service provider’s IVR system, there will be no match. In such cases, the system can be configured to fall back to a traditional grammar based approach by offering a pre-defined set of alternatives to the caller such as “I’m sorry, please let me know whether you would like billing, new plans ..” . Keyword spotting can then be applied to the response given by the customer to enable a successful interaction.

This combination of statistical analysis and keyword spotting has been very effective in contact centers for call routing applications. When callers are allowed to use their own language, complex menu navigations can be avoided and misrouting can be minimized.

Vendors offering NLP technique based solutions often offer a conversational approach based IVR systems that continuously adapt based on the user's response to previous queries. This way the caller is given a personalized service and the conversation is controlled by the caller. One of the key elements of a successful NLP based system is the ability to interpret additional information provided by the caller. If the system is able to correctly respond to this additional information, then the call duration can be significantly reduced and the call can be more productive overall. Callers often provide more information than prompted for and an intelligent system that can parse this additional information, without having to ask for it again can make a much better system than one which cannot. For example, a ticket booking IVR system may ask "Where do you want to travel" and the caller might reply "Two tickets to Boston and back for the 25th of September and flying back a week later". If the system is able to interpret the travel dates as well from this conversation, then it will definitely make the conversation more smooth and fast. It is also important to reduce the number of confirmations that the system asks as a constant stream of confirmation will make the conversation feel artificial, thus causing customer satisfaction levels to dip. Good NLP based IVR systems dynamically embed confirmations in the next prompt and handle corrections and verifications more subtly there by making the interaction more fruitful.

In short a good NLP based system must provide callers the flexibility to provide responses the way they want and the system should still be able to interpret it correctly and also able to handle multiple pieces of information in the conversation. The system should also be able to handle confirmations, without having to explicitly ask for it.

4. History and Evolution

Research in natural language processing dates back to the late forties when Machine Translation projects were initiated based on the new found understanding of parsing languages obtained during the war years by breaking down enemy codes. In 1949, Weaver brought out a memorandum which suggested the use of techniques from cryptography and information theory for language translation using computers. Soon, research institutions across the US started work in various NLP techniques. Initially researchers in Machine Translation took the simplistic approach of assuming that languages differ only in the vocabulary and word order patterns. Thus, early systems used basic dictionary look-up for word translation and then reordered the words to match the rules of the target language. Thus the contextual meaning which was inherent in natural languages was not captured by these systems leading to poor results. Soon, researchers realized that the task at hand required a more robust language theory.

In 1957 Chomsky published Syntactic Structures which was instrumental in introducing the idea of generative grammar. Soon, researchers started exploring on how mainstream linguistics could contribute to the field of NLP and MT. During the same period, other NLP applications such as speech recognition began to emerge. However, language processing and speech continued as distinct fields of research with the former focusing on theoretical aspects such as generative grammars and the latter focused on statistical methods and ignoring linguistic theories. In the fifties researchers were of the view that with the developments in parsing algorithms and linguistic theory, it would be easy to develop translation systems that match the capabilities of a human translator. However, this was not possible due to inadequacies in computer systems of the time as well as the limited linguistic knowledge.

The ALPAC report (Automatic Language Processing Advisory Committee of the National Academy of Science - National Research Council) report of 1966 concluded that MT was not an achievable target and this led to a major decrease in NLP research. In the 60s and 70s there was a lot of theoretical work around representing the meaning of language used and on developing computationally implementable solutions as the then existing transformational generative grammars were too focused on syntactic and not semantics.

As language theory moved away from syntactic to semantic representations, there were several contributions from the research community for the theoretical side of NLP. These include case grammar developed by Fillmore, conceptual dependency theory of Schank, augmented transition networks of Woods, semantic networks of Quillian and Kay's functional grammar. Along with developments in theoretical modelling, there were many prototype systems that were being built to demonstrate the effectiveness of specific theories. The ELIZA system could replicate a conversation between a psychologist and a patient by using different permutations of the user input. Other systems such as SHRDLU, PARRY and LUNAR were all used for demonstrating different natural language processing capabilities. In the late 70s the attention of NLP research was focused on semantic issues and achieving communicative goals as well as on natural language generation. Some of the key contributors to the field of NLP include Mann, Thompson, Hobbs, Rosenschein, Polanyi , Scha and Reichman.

As critical computational resources became more affordable in the 80s, and with each field of research becoming increasingly aware of their limitations, and with the need for real world solutions, researchers started focusing on non-symbolic methods. In recent years, NLP as a field has been growing rapidly, due to the increase in the amount of electronic text, the easy availability of computational power and memory and the wide prevalence of internet. Statistical approaches have been widely used in almost all NLP applications for dealing with problems in computational linguistics such as part-of-speech identification and word sense disambiguation. Today, NLP systems are capable of dealing with general text and the ambiguity and variability associated with natural language in practical settings.

5. NLP Levels

Any system that attempts to understand a natural language, must have an in-depth knowledge about the structure of the language – its vocabulary, how the words combine to form phrases and sentences, the meaning of the words and how the words combine contextually to give meaning to sentences. In addition, pragmatic knowledge of the world around us and how humans reason based on language is necessary to interpret the text successfully.

The 'levels of language' approach (synchronic model of language) has been used to define the underlying working of an NLP system. This is similar to the way humans interpret meaning. The main levels of language processing are phonological , morphological, syntactic, semantic and pragmatic. This is different from the earlier sequential model which assumed that human language processing follows certain steps in a sequential manner. However, research has now determined that language processing is more dynamic than linear and the various levels can interact in different orders depending on the situation and information from a higher level of processing can assist in a lower level analysis. For example, the pragmatic knowledge that the document that is being interpreted relates to contact centres can be used to identify the right meaning of a particular word that can have multiple meanings.

5.1. Phonology

The phonological level deals with interpretation of sounds within and across words. There are 3 main categories of phonological rules: 1. Phonetic rules for sounds within words 2. Phonemic rules for pronunciation variations when words are spoken together and 3. Prosodic rules for changes in intonation and stress across a sentence. In an NLP system that works with spoken language, the sound waves are digitized and then interpreted using various rules or by comparing it against a predefined language model.

5. 2. Morphology

This level deals with morphemes – which are the smallest meaningful units in a word. The meaning of each morpheme remains the same across words. An NLP system can thus breakdown an unknown word into morphemes and interpret its meaning. For example, adding a prefix 'pre' to a word, conveys the meaning 'before' and this can be a morphological rule in the NLP system.

5. 3. Lexical

At the lexical level the meaning of individual words are interpreted. The first step in lexical processing is to assign a part-of-speech tag to each word. Words which can function as multiple part of speech words, are tagged to the most probable part-of-speech based on the context. At this level, all words which have only one possible meaning are replaced by a semantic representation of the word meaning. The way in which the semantic representation is done may vary based on the semantic model used in the NLP system. One such model is the logical predicates model where a single lexical unit is broken down into more basic properties. A set of semantic primitives that are used across all words and complex interpretations of meaning are derived from these simplified lexical representations. Some NLP systems use a lexicon at the lexical level. The level of complexity of the lexicon is also system dependent – some systems may only store the word and its part of speech while others may contain information such as the semantic class of the word, what arguments it takes, the semantic limitations of the arguments and so on.

5. 4 Syntactic

In the syntactic level the grammatical structure (syntax) of the sentence is identified by analysing the words in the sentence. For this, the NLP system requires a grammar database and a parser application which are used to identify the structural dependency relationships among the words in a sentence. Depending on the nature of the application, either a full parsing of the sentence is performed or only phrasal and clausal dependencies are parsed. In the latter case, challenges around prepositional phrase attachment and conjunction scoping are no longer relevant. In most languages, the meaning of the sentence can be deciphered through syntax itself, as the order and dependency contribute significantly to the meaning being conveyed. For example, “The dog chased the boy” and “The boy chased the dog”, have the same set of words, but convey different meanings through just a different ordering of the words.

5.5. Semantic

While all levels of NLP processing contribute towards deciphering the meaning of a sentence, it is the semantic level where semantic disambiguation of words with multiple possible meanings is performed. This is similar to the syntactic disambiguation (performed at the syntactic level) of words which can function as different parts of speech.

This level also focuses on determining the possible meanings of a sentence by analysing the interactions between words in a sentence. Semantic disambiguation selects one meaning of the word to be selected and included in the semantic representation of the sentence. In the semantic level, information from the rest of the sentence is used for disambiguation of words and this is how the level differs from the lexical level where words are analysed in isolation. There are several algorithms that are used for semantic disambiguation. Information such as the frequency with which a particular meaning of the word occurs in relation to the rest of the words in the sentence, the frequency with which a particular meaning occurs in general usage, the domain and context of the text being analyzed etc all form key parameters in the disambiguation exercise.

5.6. Discourse

While the syntactic and semantic levels analyze sentences, the discourse level works with units of text which are larger than a sentence. In this level the properties of the text as a whole is analyzed to identify the meaning being conveyed through the connections between component sentences. There are multiple methods of discourse processing. The most common methods are anaphora resolution and discourse/text structure recognition. Anaphora resolution focuses on replacing pronouns and similar semantically vacant words with the corresponding entity. Thus, if a text says “John came home. He was tired”, the ‘he’ will be replaced by ‘John’ in this level of processing. In text structure recognition the functions of sentences in the text is classified into discourse components so as to better interpret the meaning of the text. Thus, an essay can be deconstructed into components such as Introduction, Main Body, Conclusion and References.

5.7. Pragmatic

This level of analysis focuses on the situational use of language and thus the context of usage is more relevant for this level than the content of the text being analyzed. This level is concerned with explaining the implied meaning in a text which may not be verbally included in it. For this analysis, it is necessary for the NLP system to have world knowledge, including the intentions, plans and goals of the given context. Knowledge bases and inferencing modules are used to achieve pragmatic processing of text. For example, the following sentences require resolution of the anaphoric term 'it', but this resolution requires pragmatic or world knowledge to be embedded in the NLP system.

The cat ate the food as it was hungry.

The cat ate the food as it was tasty.

Most NLP systems tend to focus only on the lower levels of processing as very few applications require the interpretation at higher levels. Another reason is that the lower levels have been more thoroughly researched and effective solutions are available at these levels. Lower levels also deal with smaller units of analysis such as morphemes, words or at most sentences. These levels can mostly be interpreted using a rule base. However, higher levels of language processing involves regularity governed texts and extensive world knowledge to be maintained. From an approach perspective statistical approaches have been used for lower level analysis whereas symbolic approaches are more useful for higher levels.

6. Technology

There are several approaches towards natural language processing. These can be broadly divided into 4 categories: symbolic, statistical, connectionist, and hybrid. The early approaches towards NLP were mostly symbolic and statistical methods with the former being more popular. However, by the 80s, statistical methods started getting more focus as computational power was readily available and the approach involved dealing with real world contexts and semantics of language.

6.1. Symbolic Approach

Symbolic approaches analyse linguistic phenomena using knowledge schemes that represent well known facts about language and associated algorithms. The primary elements in a symbolic system are based on human developed rules and lexicons of a language. Most symbolic systems have a rule based approach or a logic based approach. A rule based system consists of the following: a.) rule database which captures the knowledge of the system as a set of rules, b.)an inference engine which applies these rules on the given text and performs corresponding actions when a rule is matching and c.) a workspace or a working memory. In logic based system, the symbolic structure of the language is represented as a set of logical propositions. Manipulations of the language structure are defined by inference procedures.

Semantic networks proposed initially by Quillian are another example of a symbolic approach. Semantic networks are modelled using a set of nodes that represent knowledge concepts and labelled links that represent the relationship between the nodes. The connectivity pattern is based on the semantic association and highly related concepts are directly linked whereas weak relations are depicted through indirect linkages via intervening nodes.

Other symbolic approach techniques include inductive logic programming, K nearest neighbour algorithms, decision trees and explanation-based learning.

Symbolic approaches have been used in a wide variety of NLP applications such as information extraction, ambiguity resolution and text categorisation.

6.2. Statistical Approach

As the name suggests statistical approaches makes use of statistical and mathematical models for NLP. Large amounts of text are used to develop generalized mathematical models of linguistic phenomena. Unlike symbolic approaches, statistical approaches use actual text and not rule bases for pattern identification. Thus, statistical approaches typically do not use outside world knowledge or linguistic rules for the analysis and relies solely on actual text as the primary source of evidence.

One example of a widely used statistical approach in NLP is the Hidden Markov Model (HMM) where a set of states are defined along the probability of transitioning from one state to the other. The term 'hidden' refers to the fact that the states are kept hidden from the end user and only the output of each state is visible.

Statistical approaches have been used for various practical applications such as speech recognition, part-of-speech tagging, statistical grammar and language parsing.

6.3. Connectionist Approach

Connectionist approaches are similar to statistical approaches in that they also develop generalized models from specific examples of linguistic phenomena. However, they differ from statistical approaches in that they combine statistical learning with language representation theories, thus permitting transformation, inference and manipulation based on a rule base or logic. Further, linguistic modeling is more difficult in a connectionist approach as they are typically less constrained than a statistical architecture. In simple terms, a connectionist model can be envisaged as a network of interconnected processing units with knowledge stored as weights of the connections between the units. Local interactions between the units lead to dynamic global behavior, and resulting computations.

Connectionist models can be either localist models or distributed models. In a localist model each unit represents a specific concept. For example, one unit might represent the concept of "animal", another unit might represent "mammal" and a third unit might represent "cat". Relationships among these three concepts are modelled as weights of connections between these concepts. The language knowledge is spread across the network of units and the connectivity among the units reflect the structural relationships. These models are similar to semantic networks and the key difference is that the connections between units are not labelled in a localist model unlike in a semantic network. Localist connectionist models find practical applications in tasks such as language generation and word-sense disambiguation.

In distributed models, a concept is represented not as a single unit but as a simultaneous combination of multiple units. Thus a single unit only forms a part of a concept representation. These models are most suited for tasks such as limited domain translations and retrievals as well as syntactic parsing.

7. Practical Applications

The theory of Natural Language Processing can be applied to a host of real world scenarios. Any application that uses language in any form (text or speech) is a good candidate for NLP. Some of the key practical applications for NLP techniques are listed below:

Machine Translation: Machine translation is one of the oldest applications of NLP and various levels of NLP have been used for MT ranging from phonological level approaches to discourse level approaches to make it as human-like as possible.

Information Retrieval: Information retrieval focuses on retrieving potentially relevant documents from a large database based on a user query. Mostly statistical approaches have been used for NLP based solutions in this area.

Information Extraction: This area focuses on retrieving certain key elements of information (such as name of persons, locations or other information) from large volumes of text using language recognition, tagging and extraction into a structured representation. The extracted information can then be used for various applications ranging from an IVR system to data mining.

Question-Answering: Question-answering applications are more specific than information retrieval applications and focus on providing the exact answer to a user query.

Summarization: These applications use the higher levels of NLP such as the discourse level, to provide a shorter representation of the original larger document without losing any important information.

Dialogue Systems: NLP applications of the future shall focus on dialogue systems. Today, dialog systems focus on addressing a very narrow set of scenarios and mostly find applications in smart appliances such as a refrigerator or a home theatre system and utilize only the phonetic and lexical levels. As NLP techniques become more and more sophisticated, higher levels shall be utilized to provide fully integrated dialog systems that make man-machine interaction more realistic and effective.

7.1. Advantages of NLP Based Applications

Contact centers which use NLP based applications can obtain the capability of providing more targeted answers to customer queries without agents having to manually search for the relevant answers. This results in increased customer satisfaction, increased first call resolution rates and improves agent productivity. NLP based applications can analyze caller queries to identify the true intent of each query and provide near real-time contextual data from the underlying databases to quickly and efficiently respond to the caller with the most-relevant answers. Some of the key advantages of using NLP based systems in a contact center are:

Minimize agent search time: With NLP based solutions which use industry domain dictionaries along with other rule bases, it is possible to automate searching for responses to customer queries more accurate, thus minimizing the agent time spent on this.

Improved collaboration: With NLP based automation freeing agent time, there are more opportunities for agents to collaborate with each other, share ideas and expertise.

Enhanced customer satisfaction: NLP based applications can dynamically link multiple applications and knowledge databases to enable improvement of key call center metrics such as an increase in First Call Resolutions and reduced call durations, thus leading to enhanced customer satisfaction.

8. Conclusion

NLP is still a very nascent area of research. However, there have been sufficient advances in NLP technology and practical applications to believe that NLP is here to stay. As economic conditions become difficult enterprises are looking for ways to cut costs without compromising on service quality. Contact centres which rely on customer communications as the basis of their business model can achieve significant cost savings by adopting natural language techniques based self-service IVR systems to increase automation rates. There are different approaches of NLP that can be applied at different stages of the customer interaction to achieve improvements in automation, reduction in operational costs and increase in customer satisfaction levels. With the help of NLP techniques callers can be provided the flexibility to give responses in a natural way to the IVR system and this can lead to more efficient and pleasant interactions with the contact center.

As the bar for service is being raised across industries, contact centers are expected to provide customers with intelligent and accurate answers as fast as possible and still keeping operational costs low. As contact centers keep looking out for ways in which more can be achieved with less, automation using cutting-edge technology such as NLP could be the answer for delivering exceptional customer service while being able to cut down on agent staffing and budgets. This is not to say that automation shall replace agents. On the contrary, as the volume and diversity of customer enquiries keep rising and multi-channel service delivery becomes the norm rather than the exception, techniques such as NLP shall provide agents with the much needed tools and information which enable them to deliver superior customer service, thus enabling a win-win situation for the agents, the contact center as well as the customer.

References

1. http://china.nuance.com/ucmprod/groups/enterprise/@web/@enus/documents/collateral/nd_005362.pdf
2. http://surface.syr.edu/cgi/viewcontent.cgi?article=1043&context=istpub