

Bots

It's Good News for the Enterprise

Dr. Siddhartha Chatterjee

Purushottam Darshankar

Contents

- Bots in the Enterprise **4**
- Automatic Speech Recognition **5**
- Natural Language Processing **6**
- Know User Intent and Context..... **6**
- Knowledge Representation and Information Retrieval..... **8**
- Bots — Give them a Hug..... **9**



It is interesting how life imitates art, and vice versa. Recall the scene pictured here from the 1986 movie “Star Trek: The Voyage Home”, where Chief Engineer Montgomery Scott attempts (unsuccessfully) to talk to a Mac by speaking into its mouse.

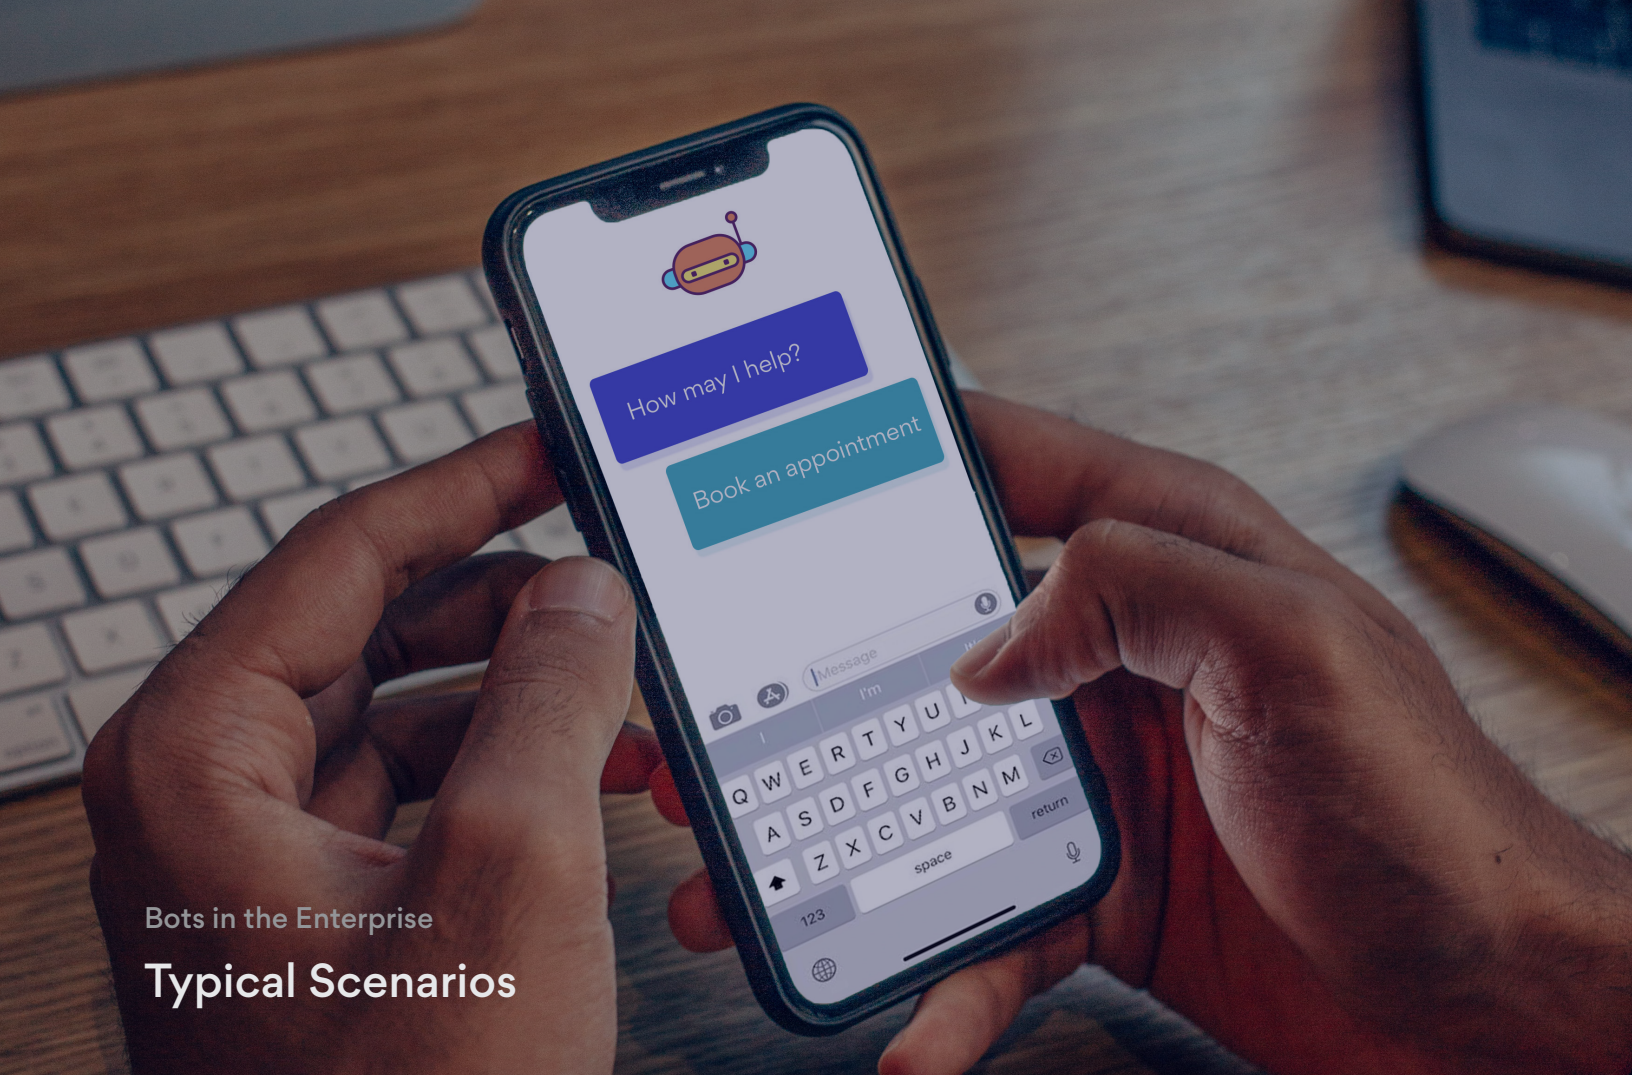
If the movie had been made today, he probably would have been shown talking to a computing device, although the prompt would have been “Siri” or “Ok Google” or “Alexa” rather than “Computer!” Bots (short for “chatbots”, or “Conversational Interfaces” to give them their proper name) are an exciting new trend — pushed by a software-driven world — primed to make our lives better. In the coming years, we expect enterprises to adopt bot platforms in the same way they are currently embracing cloud, IoT, and digital transformation platforms.

The capabilities of bots is, of course, driven by the significant growth in relevant technologies, as we will discuss below. However, enterprise interest in them is driven by two factors.

1\ Businesses are under relentless pressure to cut down costs, and bots provide a path towards this. In fact, such capabilities have been widespread in the phone world for a while,

although they tend to be transactional rather than conversational in nature, and they are often more of a nuisance than a help.

2\ Enterprises are also under pressure to deliver better experiences to their customers and create more intelligent and intuitive interactions. The days of green screens are long gone; WIMP, KVM, and GUI are acronyms from the past century; and apps are giving way to “experiences”. Enterprise clients (i.e., each one of us) who have enjoyed these technologies now for a while in the consumer domain (Apple Siri, Google Now, Microsoft Cortana, or Amazon Alexa — depending on your platform preference) are now expecting the same in their interactions with financial institutions, healthcare providers, retailers, and other enterprises.



Bots in the Enterprise

Typical Scenarios

Advances in AI technologies like Natural Language Processing (NLP), Deep Learning (DL) and Machine Learning (ML) have created intelligent bots that are set to impact every aspect of enterprise processes and workflows.

The bots are helping enterprises convert labyrinthine processes to simpler forms that are revolving around customer experiences and are being converted from system of records to system of engagements. The variety of enterprise processes and workflows that can be converted in this manner will continue to increase as businesses continue to take the lead in adopting intelligent bots.

Here are a few enterprise scenarios that users are already experiencing or will see in the near future.

When calling a customer care center or using an IM function on the help section of a company's website, the user is invariably interacting with bot agents that will run all service operations with minimal or no human intervention. On the phone, such interactions might be very obvious (from the

robotic voice, if nothing else) and we have grown accustomed to this sort of interaction. Online, where the interaction can be in a chat window in textual form, the interaction can be less obvious and it can actually feel like we are chatting with another person. Of course the rub comes when the customer has a question that the bot is unable to answer. We all mostly likely have had experiences where we needed to throw an exception and be re-routed to a human customer service representative. This emphasizes the need to have an actual human in the loop, who is monitoring dozens of ongoing bot conversations, and can step into the conversation when the bot is unable to provide the customer with a relevant response and appropriate experience.

Retail and travel are probably the hottest or most easily evident areas where we see this move to a software-driven world. When you go online, bots will provide you with intelligent recommendations based on your interest and purchase history. They are utilizing two aspects of actionable insights — using data to make intelligent decisions — to try and influence your purchase. The first is your history of

personal preferences and habits. The second is utilizing big data-based learning and predictive algorithms where they can analyze patterns based on millions of transactions to determine the greatest likelihood of what will appeal to you. The software running behind these bots will constantly become smarter and smarter as every digital interaction will be stored and used in building the context during conversation.

In the B2E scenario, bots let employees interact directly with enterprise applications that matter to them inside the chat client (like Lync, Yammer, Slack, Salesforce Chatter etc.) — and bring a level of automation.

As everything in the world becomes software, bots will enable field technicians to collect and present

information regarding an issue, and also as a resource, when they need help or find themselves stuck on a service assignment. Or taken one step further, bots will help with self-service by providing the customer options that lead to an immediate solution for fixing the issue. All this is in real-time, with assistance that is personalized and available instantly.

With that as our background, we now give you a look at the process we undertook to building an enterprise bot platform and provide a summary of our observations from our initial experimentation.

“Chat bots are computer programs that mimic conversation with people using artificial intelligence. They can transform the way you interact with the internet from a series of self initiated tasks to a quasi-conversation.”

As described by The Guardian

Automatic Speech Recognition

Race for an Edge in the Global Market

“Ok Google, will it rain today?” Both Apple and Google have done a good job educating users on the value and ease of voice controlled features. We have come a long way from the 1980s when we were amazed we could speak our account number into the phone! The no-UI trend represents software in its purest form: humans talking directly to machines through voice commands.

Speech is now widely accepted and replacing or complementing traditional input devices such as the mouse, keyboards, or touch sensitive screens. With advancement of machine learning, this technology has now approached viability for speech-based human machine interaction. Without question, people are now ready to embrace the simplicity of interacting with intelligent bots.

There are slew of significant players who have emerged on the scene, including Google, Apple, Amazon and Microsoft. The following criteria was considered for evaluation.

- \ Unified and cross-platform operation through availability of REST interfaces
- \ Quality of speech recognition, with focus on lexical, syntactic, and semantic aspects of language
- \ Speech recognition provided as a service and the associated cost models

Google Cloud Speech APIs are not yet production-ready and do not come with any official documentation and usage capabilities.

It does, however, provide Speech API as a native feature on Android smartphone devices.

Nuance provides REST APIs with a registration key and does not come free for commercial rollout. Baidu through its DeepSpeech 2 technology is claiming a breakthrough solution — China's answer to Google. But the jury is still out.

For our evaluation, we used different phrases and

different gender and background noises.

Here is a quick observation summary — when tested from desktop browser, Android, and iOS devices, Google is by far the best out-of-the-box solution for Android devices followed by Nuance, which provides cross-platform APIs that can be seamlessly integrated into bot solution. Other solutions from Watson, api.ai, and Baidu were evaluated but they were not close to the performance that was delivered by either Google or Nuance.

Natural Language Processing

Conversations are Hard to Automate

Understanding spoken language and free text conversation is not as straightforward as it might seem at first look. Actually, it may not seem straightforward at all. Semantically, the same command can be expressed using different sentences, synonymous phrases, etc. Providing a rich natural language interface is essential to provide a strong user experience with bots.

In order to address those challenges, enterprise bot solutions include natural language processing (NLP)

systems that can identify phrases and synonymous phrases. Some of the commonly associated tasks for NLP are text segmentation, names-entity recognition, stemming, and stop word removal.

There are plenty of tools available for NLP processing, the most popular being Stanford NLP, NLTK, Watson and GATE. All of them perform well, and you should select one based on how easily it fits into your overall technology stack.

Know User Intent and Context

Look beyond Spoken Language Processing

In order to build good conversational interfaces we need to look beyond the simple keyword search by substring or regular expression that we usually use.

Bots need to understand what the user wants to accomplish and collect a bunch of information about the task in order to perform an action. And this needs to be done in real-time. At every interaction, based on the history and the current user input, the bot should either request information, ask for confirmation, or perform some kind of action.

Thus intent is the core concept in building the conversational bot. So in order to understand user intent, we need to map a phrase to a specific action that we can really provide. Along with intent, it is necessary to extract the entities or parameters to perform the intended task. Examples of entities

types that are commonly used include — locations, date/time, numbers, organization names, duration, currency, and enumerations of options.

The intermediate states and parameters from previous utterances during the conversation need to be stored in order to build the context. When the bot interacts with a user, it needs to manage the state and change the actions based on context. All of this needs to be implemented by the developer, often using the exposed APIs by the bot framework. There are many players in the market who provide APIs for extraction of entities, intent, and sentiments with some support for managing the conversation.



“Nearly 3 out of 4 consumers prefer to solve their customer service issues on their own, setting the stage for big changes in the consumer experience landscape.”

Study by ASPECT 2015

IBM Watson provides a set of text understanding APIs under [the Watson Developer Cloud and AlchemyAPI](#). There is no single API that does intent and entity recognition in a single call. You need multiple API calls to extract all the information you need, and it's harder to link entities to intent.

Microsoft provides the [Language Understanding Intelligent Service \(LUIS\)](#), which has an API that can do both intention and entity recognition at the same time. You can bind an action by calling another API to a set of intent and entities. It also has a list of pre-built intents and entities.

[Wit.ai](#), acquired by Facebook in 2015, uses a bot engine that learns conversation flow from examples of user input plus bot response. As it does not provide action support, calling external services must be implemented outside the platform.

[Api.ai](#) now acquired by Google is also similar to [wit.ai](#) with several pre-trained intents already made available.

[Amazon Alexa Skills Kit](#) is another good system that you should keep in mind. It is currently a leading household platform for conversations which Amazon is aggressively pushing forward with new offerings and features.

Most of the APIs are good to get started but you may experience limitations. To solve real-world problems, developers will be pressed to look for external tools to provide more sophistication and make them work smoothly in an enterprise environment.

Keys to Success

The success of a bot is measured based on its knowledge representation ability as well as its information retrieval and dialog capabilities. During the conversation, the bot looks for matches between the user query and set of answer modules stored in its knowledge base.

One of the approaches for knowledge representation is ontology. In philosophy, ontology is the metaphysical study of the nature of things. In computer science, ontology is a formal naming and definition of the types, properties and interrelationships of the entities that really or fundamentally exist for a particular domain of discourse. In other words, it is a practical application of philosophical ontology with a taxonomy.

In this structure, resources are compiled by extracting certain kinds of information from previously designated textual resources. It uses Resource Description Framework (RDF) triples that are encoded in an XML format, in order to represent machine processable concepts. **YAGO** (Yet Another Great Ontology) and **Wordnet** are the good examples of most commonly used ontologies.

Machine learning provides the means to building effective information retrieval systems that have access to collections of thousands, or millions, of documents, from which, by providing an appropriate description, users can recover any one. Typically, users iteratively refine the descriptions they provide to satisfy their needs, and retrieval systems can utilize user feedback on selected documents to indicate the accuracy of the description at any stage.

The most popular representation is based on the “bag of words” or term vector model, where a document is represented by a set of terms. Documents that contain a given set of terms each correspond to a certain point in n -dimensional space, where n is the number of terms. A term vector represents a set of words that appear in the document whose values are derived from the number of occurrences in a document — the more often a term is mentioned, the more likely it is to be central to the subject matter. The issue is, however, that there are different ways to refer to the same object or concept, which do not necessarily have words in common — the problem of synonymy. Conversely, many words mean quite different things in different contexts — the problem of polysemy.

In order to circumvent these problem, other methods such as latent semantic analysis (LSA) are used, in which term vectors are remapped to a different semantic space to support better conversational capabilities.

Recent years have seen significant progress in information retrieval and natural language processing with deep learning technologies. The state of the art solution is to use a neural network that can help in building open domain, conversational dialog systems based on large dialog corpora using generative models. Generative models produce system responses that are autonomously generated word-by-word, opening up the possibility for realistic, flexible interactions.

25% of homes will use smart assistants to access different online services by 2019.

As predicted by Gartner

Bots

Give them a Hug

The goal of bots is to make customers feel as if they are communicating with another human, rather than a piece of intelligent software, in an environment that calls for little to no human intervention (in most cases when people prefer that). The goal is also to reduce costs for businesses in manning such areas as customer service. Gartner sees 33% of all customer service interactions as still needing a human intermediary by 2017, down from nearly 60% in 2014.

With the growing interest of organizations to have a closer, real-time connection with their customers, it makes great sense for businesses to consider

B2C (business-to-consumer) scenarios for bot automation, take advantage of early adoption and beat competition as customers feel more valued.

However, when it comes to transitioning from existing technologies to emerging technologies, such as bots, one needs to evaluate how the technology, its ecosystem, and customer needs might evolve and prepare for new alternatives.

While we will have to wait a little longer to see how the bots wave plays out over the next few months, early responses and clear benefits indicates that it is time for businesses to embrace them.

About Persistent

Persistent Systems (BSE & NSE: PERSISTENT) builds software that drives our customers' business; enterprises and software product companies with software at the core of their digital transformation.

www.persistent.com

India

Persistent Systems Limited
Bhageerath, 402,
Senapati Bapat Road
Pune 411016.
Tel: +91 (20) 6703 0000
Fax: +91 (20) 6703 0008

USA

Persistent Systems, Inc.
2055 Laurelwood Road, Suite 210
Santa Clara, CA 95054
Tel: +1 (408) 216 7010
Fax: +1 (408) 451 9177
Email: info@persistent.com



Persistent