REAL AI

A Definitive Handbook

By Navveen Balani

REAL AI

A Definitive Handbook

By Navveen Balani

Copyright @ 2019 by Navveen Balani. (http://navveenbalani.com)

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, printing, recording, or other electronic or mechanical methods, without the prior written permission of the publisher.

All other trademarks or registered trademarks are the property of their respective owners.

March 2019: First Draft

Although the author and publisher have made every effort to ensure that the information in this book was correct at press time, the author and publisher do not assume and hereby disclaim any liability to any party for any loss, damage, or disruption caused by errors or omissions, whether such errors or omissions result from negligence, accident, or any other cause. Use of the information and instructions contained in this work is at your own risk.

Introduction

Since time immemorial, the human race has been fascinated with what machines could or couldn't do. The quest for creating intelligence that can surpass human intelligence continues to grow and its coined in a term called "Artificial Intelligence (AI)".

The term AI though was coined by Mr John McCarthy in 1955, the computer visionary who I am sure would have been glad that his Shakespearean display caught the eye of researchers around the world as it ranked 4th in the Scopus scholarly database. It's even fascinating to know that it didn't appear in the top 10 searches the year before. A spectacular performance. One could relate it to a New blockbuster film opening weekend where it overtakes fading glories. But in the case of AI, it would be fair to say it's a sequel story. The sequel story is only behind cancer for obvious reasons. (As most researchers/scientists belong to the medical discipline and the cure to cancer remains the last undiscovered El dorado in the field). The other 2 words in the list are two rather newly coined terms Blockchain and Big Data. No surprises here if you are following the technology trends, as all the 3 are linked to AI in some way or another, with a potential to create intelligent and trusted systems in future.

AI is the latest marketing buzzword that is made to find its place in every possible use case - from driverless cars to intelligent chatbots, from robots like Sophia to solving cancer problems, from winning games to providing human like intelligence. But is this current hype real or we have just started scratching the surface of intelligence.

In order to make the distinction on hype v/s reality, let's go in some basic technical details of AI technology.

AI stands for artificial intelligence. It’s an intelligence system put together artificially to learn and provide an output. Learning can be done by providing data to the AI system. Data can be big data, customer data, unstructured text, audio, visual, environment surrounding details etc. Based on the data provided, an AI system would learn and identify hidden patterns and provide an output.

For instance, if an AI is recommending what food to order, it must know your food preferences, what you had ordered before, where do you usually order from, what days you usually order specific cuisine and lot of other details to recommend the right cuisine for you. The output can be a list of Top 5 food orders for today.

Similarly, if an AI is assisting a doctor for providing options for cancer treatment, the system must have the complete patient medical history, must understand the complete cancer domain (or the respective specialization) and also periodically learn any new treatments or findings from medical journals. Understanding the complete cancer domain is a very complex process, where one needs to train the system to understand the medical terminology and the vast ever-growing cancer literature, identify patterns and correlations from existing patients, their suggestive treatments and outcome and finally suggest options for treatment. There can be many more data points and this is a continuous process where system would be trained from the feedback and their outcome. While we keep hearing, AI is helping solve cancer cases, this is far from reality and systems have just started to touch the surface.

To make life simpler, just remember the following distinction -

"**AI can learn, but can't think**".

Thinking would always be left to Human on how to use the output of an AI system. AI systems and their knowledge would always be boxed to what it has learned, but can never be generalized (like humans) to

think outside the domain it has been trained on. Understanding this distinction is very important. Human intelligence with only few sets of observations can learn, think and apply their learnings on different domains quite easily. A simple example would be of a doctor treating cancer patient can give you advice for common cold, but an AI system trained specifically on cancer data, may not even understand what common cold means, leave aside the treatment options. Building a generalized AI system may or may not happen in future. The current focus should be building domain specific intelligence and get it right.

AI can never be a replacement for Human Intelligence. While simple to medium outputs of AI can be automated to skip a Human expert, the majority of the decision making and critical intelligence would always need Human intelligence.

While AI is being projected as the next big technology that can transform our world, we are far from away in releasing this vision. You may hear many successful AI marketing strategies, but AI is yet to deliver its true value. AI alone will not lead to transformation, but a combinatorial power of various technologies and advancements in computing power would bring it closer to its true potential.

Through this book, I plan to provide a realistic view on what AI system can achieve in today's environment and what to expect in future. I plan to draw the reality and bring you closer to Real-AI. Hence, the book is titled - "Real AI".

After going through several iterations on the format of the book, I started writing this book in a Question and Answer format as it provides direct answers to some of the questions the readers would want to know. The book is being written and available on my website - **http://navveenbalani.com/index.php/books/download-real-ai-a-definitive-handbook/**.

The book will cover the following topics -

- AI Introduction - Real facts minus the hype.
- AI chatbots - The not so intelligent chatbots.
- Recommenders - The race towards personalized recommendations
- Predictions - The illusion of AI surpassing human intelligence
- Computational Creativity - An AI that can paint, sing or dance
- What's in the future - More buzzwords to keep you busy - The Ethical AI, Explainable AI, Auditable AI and the list would go on.

One chapter of the book (AI chatbots - The not so intelligent chatbots.) is available now.

The Real AI book is part of our "The Definitive handbook" series. Our vision in the – "The Definitive handbook" series is to enable our readers to understand the technology in simple terms and provide a practical go-to reference and a recipe for building any real-world application using the latest technology.

If you have any questions and comments on the book, please write to me at me@navveenbalani.com

The opinions and views that I have provided in my book are my own.

Table of Contents

Chapter 1. AI CHATBOTS

The not so Intelligent chatbots

Welcome to the world of intelligent chatbots, your companion and conversation agents which would make your life smarter. A leading research paper even said by 2020, the average person will have more conversations with bots than with their spouse. So be ready to embrace this new life in a year from now.

Ok hold on, have you ever tried telling Siri or Google to “Find restaurants which doesn’t serve pizza”. At least they are both consistent in some way, they gave the same answer - suggesting restaurants which serve pizza.

Ok how about Sofia, the first citizen humanoid robot, which is making its way to every media event and giving interviews and boost of human like conversations. Well the truth is far from reality, it is providing an illusion of understanding conversation, but as you start asking intelligent questions you would realize it can answer fixed set of questions.

Well by now, you should be able to clear out the noise from reality. So, should I invest in chatbots with all these limitations? Yes, any technology would have its limitations, but you need to be aware of what you can build now, what to avoid and how to work around the limitations.

I have seen many companies trying to build sophisticated chatbots using products from leading chatbot vendors and cloud offerings, spending million on dollars and hitting a roadblock. If you go by what is being projected and start building it out, you would soon realize these limitations one way or the other. The problem is that most of the vendors claim it's very easy to build a chatbot, but in reality, all of these techniques fall short when it comes to building a true conversational agent.

With current implementations of chatbot, we are probably at the first generation of AI chatbots which are more or less scripted and giving answers to pointed questions. What I mean by scripted is that it is trained to understand general vocabulary, entities, the metaphor, synonyms etc. The chatbot uses fixed set of flows to understand the context. For domain specific use cases, additional training is required, and you need to train on specific domain terminology and relationship between the words. For instance, if you are building a shopping advisor chatbot, the term "black and white dress" implies "black and white" as color and dress as category. You might expect the color "black and white" is fairly generic and should be easily identified by the AI system, but that's not really the case, which I will go through (any many such examples) during the course of the book.

In this chapter, I would like to draw your attention on how to build AI chatbots the right way, understand what implementation exists today, current limitations of the technology and how to work around it. After reading this chapter, hopefully you have learned something new and can take informed discussions of building successful AI chatbots.

What are Chatbots?

A chatbot is a software program which carries out a conversation with a human. The conversation can be through textual methods, voice or even through recognizing human expressions.
Chatbot interactions can range from simple questions being answered like - “what is the outside temperature’, to sophisticated use cases which requires a series of dialogue to arrive at an outcome - like a chatbot for booking holiday trips or providing financial advice.

What are the technologies used to build Chatbots?

Chatbots are not a new concept. Earlier technologies using fixed set of input from user to drive conversations or scanned the input message to find keywords and lookup information/responses from database. These were mostly rules based and keyword driven, without understanding the context and meaning of the input message. Based on the input, a predefined programmed response would be provided.

With the advent of AI, Chatbots uses technologies like Natural Language processing to understand the language and intent from the input message and take corrective action. As the system tries to understand the language, users asking the same questions in multiple ways, the system is now able to understand the intent. Once the intent is identified, you can extract the interested topic from the input.

Info - Natural language processing (NLP) is a branch of AI to help systems understand, interpret and process human languages.

For instance -
Find the cheapest flight from US to UK is similar to Find me lowest air fare from US to UK.
Here the intent is - cheapest or lowest flight
Topics are - Location: From location - US, To Location UK
Action - Search flights.

An AI open source package or an AI NLP cloud service can be used to develop chatbots. Let's refer to this as chatbot implementation for future references. We would talk about chatbot implementation in detail during the course of this chapter

What should I keep in mind for developing an AI Chatbot?

Chatbots work well when domain is well understood by the AI system. As the AI chatbot relies on NLP to understand the semantics of the input message, unless the NLP parser is trained on the domain, the accuracy of recognizing the intent and topics of interest would be very low or not as per acceptable criteria.

Take an example of a shopping chatbot which advises user what to buy based on the latest fashion trends.
Consider 3 queries below from a user -
Query 1 - Show me medium size trending black and white dresses for Christmas party
Query 2 - Show me white color, 3 inches platform heels

Query 3 - Find And black and white floral dress under 2000

Here the chatbot needs to understand the following

- Understand the shopping language.
- Understand the intent - It's a shopping query
- Understand the domain - Its shopping query for apparel and shoes. (i.e. there can be multiple domains - grocery, electronics, books etc.)
- Understand clothing shopping category and terminology -
    - Category - dresses, sandals etc.
    - Variants - sizes (medium/large etc.), color (various colors and combinations like black and white), heel size (3 inches. etc.)
    - Prices and ranges - 2000, etc.
    - Brands like - AND, Nike etc.

Out of the box, any chatbot implementation wouldn't understand the domain. You need to train the chatbot on the custom domain to recognize the context and the language.

For instance, out of the box NLP parsers would not recognize "AND" as a brand. Let's inspect how well some of the leading Cloud AI NLP services recognizes the sentence - "Find And black and white floral dress under 2000"

Here is a snapshot from Watson NLP (out of the box) implementation.

Figure: Keywords from Watson NLP

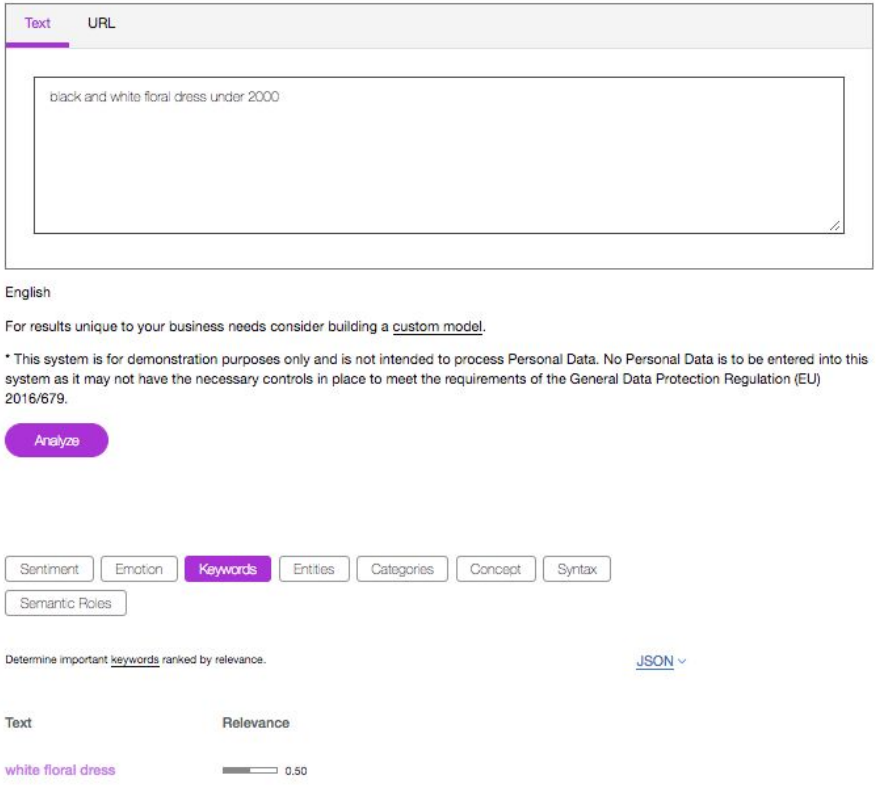

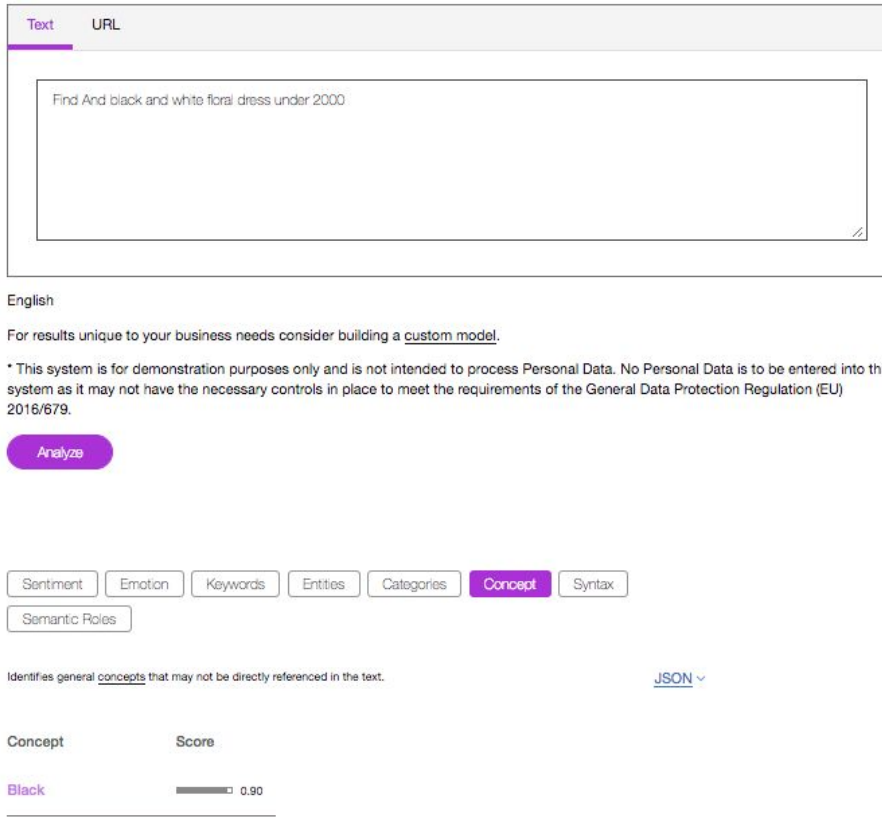

Figure: Concepts from Watson NLP

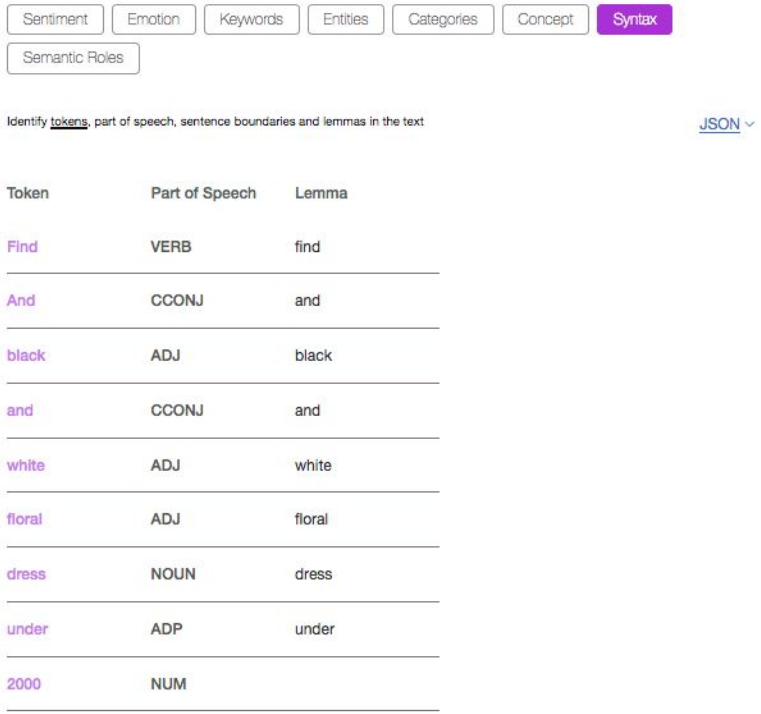

Figure: Part of Speech from Watson NLP

As you see, the Watson NLP recognizes “white floral dress” as keywords and “Black” as concept. Ideally it should have recognized “black and white” as concept as we are looking for a combination of these colors. The dress could also be a concept, as its quite generic. The floral can be keyword which has a dependency on dress. Identifying all the facts in the right way it's important, as based on the facts you would convert this to a search query to get the required details from the data store (or from respective search indexes).

For instance, the above should result as -

Color = “black and white”
Category = “Dress”

Gender = “Female”
Price < 2000
Pattern = “floral” or Keyword within category = “floral”

(where color, category, gender, price, pattern are all the columns or indexes you are searching against.)

The Watson NLP parser doesn’t recognize “And” as brand and recognizes “And” as a conjunction (“CCONJ”) in part of speech, which is expected as its not trained on it.

Let’s check how Google NP classifies this sentence. Here is a snapshot from Google NLP.

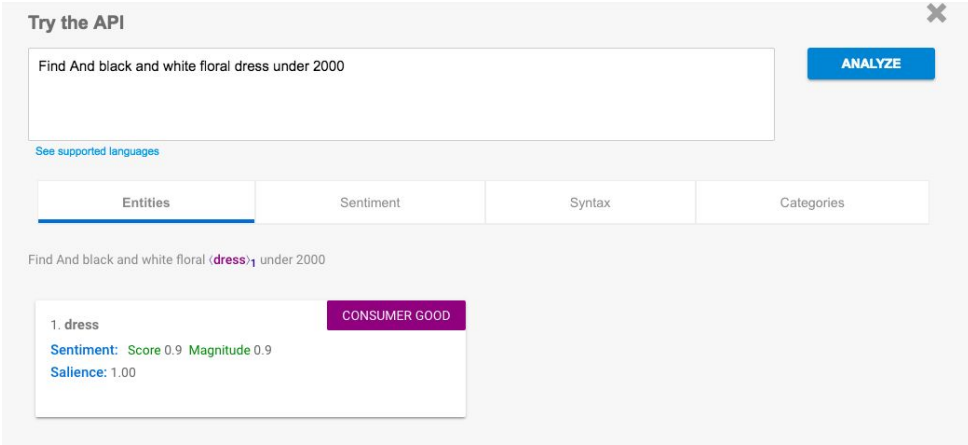

Figure: Entity classification from Google NLP

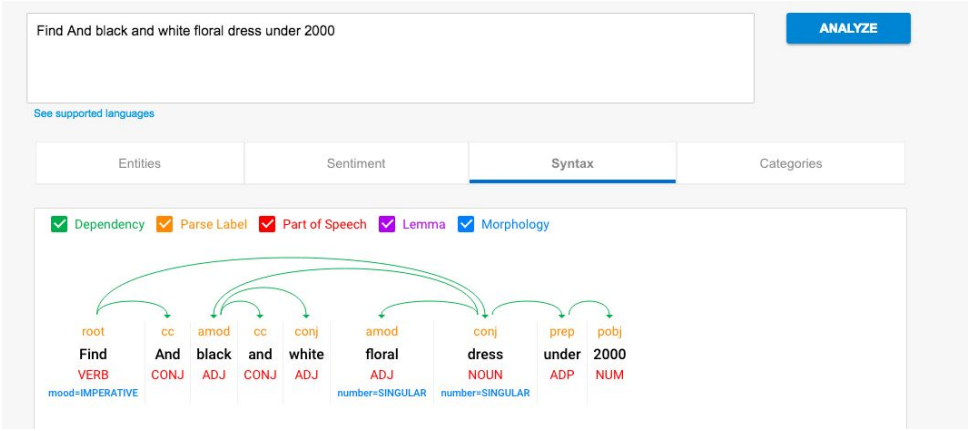

Figure: Part of Speech from Google NLP

As you see from above figures, Google NLP identify the entity as “dress”, but doesn’t identify the colors “black and white”. With respect to part of speed tagging, its similar to Watson NLP, recognizing “And” and as conjunction (“CONJ”) and not as brand.

The above is true for any of the available NLP implementation (that is available today), where it fails to understand all the correct context of the sentence. The use case was pretty simple. Even if we train the NLP implementation on these examples, it would fall short as you need to plugin specific NLP rules for such conditions to get the desired results. As the complexity and context that needs to be inferred increases, training would also not help as you can never come up with a generalized model for such conditions. That is the single most limitation if we only rely on today’s generation of NLP implementation.

Tip - Based on my experience on building a sophisticated shopping personalized advisor, none of the out of the box AI NLP implementation fitted the requirements. A simple scenario is these 3 sets of sentences - “black and white dress”, “and black dress” and “blue jeans and white shirt”. In all these 3 examples, the use of word “and” means different meaning. In the first case, its represents a combined color ““black and white”, in second instance “and” represent a brand and in third instance two queries joined by a conjunction (i.e. and). Even with required training, a generalizing model was not possible with any of the available solutions. These are just one of the many examples I am highlighting. Imagine the complexity when dealing with medical literature. In our case, we ended up building our domain specific NLP implementation which worked for all such scenarios.

In general, while designing chatbot solutions, start with a closed domain and what kind of questions the chatbot needs to answer. Don’t start building a general purpose chatbot from start, as it would be difficult to get the required accuracy. Secondly, if you are using any cloud vendor or third party implementation, ensure your use cases can be simply solved by the default implementation or you need to build components to work around it.

What are typical use cases for building a chatbot?

In today’s digital age, customers are looking for instant information and speedy resolution to all their queries.
Chatbots provides an efficient way to stay connected with end customers directly and provides information at their fingertips - be it through a messaging chat application or through voice enabled service like Alexa and Google Home.

Some of the typical use cases are listed below -

- Ability to know your customers and directly interact with them over various channels, like retail brands directly connecting to their end customers.
- Improve customer engagement, interaction and provide speedy resolution.
- Scaling customer service operations by providing relevant information 24 by 7 at customer's finger trip.
- Understand customers and their preferences better to provide hyper personalized service, like a personal assistant.
- Provide an ability to interact with connected devices, like Smart Homes in a natural and intuitive way.
- Provide expert guidance, like a financial assistant chatbot providing investment suggestions.

What are the high level steps for building an AI chatbot?

Following are the high level steps to build an AI chatbot

- Define the business use case and end goal for building the chatbot.

- Define Conversation interfaces
  - Define what kind of questions needs to be answered
  - Define conversation/dialog flow on how various interaction would happen with the user. For instance, booking a flight is one dialog flow, booking a hotel is another dialog flow, etc. Within a dialog flow, what would be interaction flow with the user.

    - Define how to capture the feedback from the user regarding the answers provided. Feedback can be explicit, like the user rating the answer or implicit on how much time a user spends looking at the answer and follow up activity after that.

- Question / Answer exploration
    - Identify existing sources (if any) for questions, like website FAQ, call center logs etc.
    - Create representation of Questions that would be asked.
    - Create variations of questions for training the chatbot to understand the language and be able to generalize well.
    - Identify source of answers - whether it would be programmed response or coming from internal knowledge sources and documents (like available technical manuals for troubleshooting device related queries)

- Pick up a Technology approach

In this step, you will decide how to implement the chatbot. There are 2 approaches - building your own chatbot implementation using available frameworks (like TensorFlow, NLP implementations like NLPTK) and custom components or using an existing platform service like Google NLP, Amazon Lex or Azure Bot service.
In both the approaches, you would need the train the chatbot implementation to recognize the question intent, domain and the language. Existing platform services have simplified this process by providing required utilities that makes it easier to create chatbots. For more details, kindly refer to “How do you build chatbot using chatbot platforms”.

- Pick up a delivery channel

In this step, you will decide how to expose the

chatbot to end users through the required channel. The channel can be web, mobiles or voice enabled devices.

Your chat implementation would typically expose an API (to ask questions and get responses) which can be called by a channel implementation. You can also release your chatbot implementation over third party services like Facebook Messenger or voice enabled services like Amazon Alexa. For more details, kindly refer to “How do you Integrate chatbots with third party services”.

- Release, Monitoring and Feedback

Once the chatbot is released, you would typically store all the user interactions to help you analyze the user behavior and their preferences better. The user and behavior data in turn would be used to provide a more personalized service. How would you use this new user information, depends on your use case. For instance, if a travel chatbot is recommending a new holiday trip, it can suggest options based on your last trip interaction. You need a build a recommendation system that looks at the history of the user interaction in the past and suggest options. For details on how to build recommendation systems, kindly refer to Recommendations Chapter.

Another important point is to capture feedback from the user at regular intervals to understand if chatbot is providing the right information. The feedback captured will be used to improve the chatbot implementation, which can lead to training the chatbot implementation with new information. For instance, your chatbot may not be trained on recognizing certain entities and concepts and as a result the responses would not be proper. You need to plan for building and releasing incremental models based on the feedback.

How do you Integrate chatbots with third party services?

As part of your chatbot technology implementation, your chat implementation would typically expose an API (to ask questions and get responses) which can be called by a channel implementation

The channel can be web, mobiles or voice enabled devices. If you already have an existing mobile application, you can embed this as part of the mobile application.

You can also release your chatbot implementation through third party chat enabled services like Facebook messaging application or through voice enable service like Amazon Alexa as a skill.
All of these chat enabled services provides a framework to plugin your own implementation. The framework provides hooks or code interceptors for intercepting the chat message. You need to extend their framework and plugin your own implementation For example, if a user asks a question on Facebook messenger, the question would be handed to your chat implementation through predefined hooks. You would process the message and send the response back, which would be sent back to the user.

Similarly, if you need to make your chatbot available over Alexa, you need to wrap it as an Alexa Skill using Alexa Skills Kit interface. Once your skill is enabled in Alexa by the user, any voice messages would be intercepted by your skill and you can provide the required implementation and responses as per your chatbot.

For more details, kindly refer to “How do you build chatbot using chatbot platforms”.

How do you build chatbot using chatbot platforms?

A chatbot platform provides you a set of services to design, develop and deploy your chatbot. They provide you with a framework and guided set of utilities to build a chatbot.
Cloud providers like AWS, Azure, IBM, Google Cloud provides you a set of services that help you to create conversations, understand the conversation language using NLP techniques, hooks to take required action and deliver the solution via APIs.

The fundamental approach adopted by each of these providers is same. They allow developers to

- Design conversation flows using some visual interface or tooling provided by cloud provider
- Through these conversation flows you
    - Provide a set of questions and multiple ways you can ask the same question
    - Define what is the Intent of the question. For example, for the question -" Find cheapest flight from US to UK", the intent is to find the lowest air fare.
    - What entities of interest to extract from the Intent. The chatbot provider needs to be made aware of these entities. In above example, entities are country list -, UK, US. These entities can be generic which are recognized automatically by the cloud provider or the cloud provider provides a mechanism where you can provide or train these entities (including synonyms etc.) through some tooling provided by the cloud provider.
    - Use the entities extracted to carry out the required action for the intent. For instance, in the above example, call a flight API service providing UK and US as "from" and "to" locations.
    - Provide the response.

- Test and expose the chatbot through an endpoint
    - The cloud vendor typically provides an ability to expose the functionality for your chatbot through an endpoint, like a REST API.

The above technology work for simple to medium complexity flow - like FAQ, pointed questions and answers for customer query, fixed step of steps (booking a cab etc.) etc. Anything which requires sophisticated handling of queries, like the shopping advisor example, needs to be custom developed using NLP and other techniques.

Info - Microsoft has a QnA service (https://www.qnamaker.ai) that lets you create bot from FAQ.

What is not real about Chatbots?

Chatbot are examples of Weak AI (we discussed types of AI in Chapter 1). Current generation of chatbot can be thought of smart dialog systems driven through techniques like NLP and fixed conversation flows.

Out of the box, a chatbot doesn’t understand any domain. We need to train the chatbot to understand the domain. Also, based on the complexity of the domain, you would incrementally train and add subdomains. For instance, a chatbot helping you book a cab is an example of fixed domain, while a chatbot helping assisting doctors for cancer treatment would be trained on various types of cancer incrementally. As mentioned in the shopping advisor example, understanding the meaning of the same word in different context is difficult for the current generation of NLP implementation to resolve and you need to rely on custom techniques to handle such conditions.

Now, let's look at some marketing gimmicks around AI chatbots -

- INGEST AND KNOW IT ALL chatbots - These are chatbots being marketed where you can ingest millions of documents, like medical literature and can ask questions, which can provide expert assistance like diagnosis of diseases. Such kinds of systems unless trained appropriately will never provide desired accuracy.

By appropriately, I mean it can take years to train these systems. The fundamental problem with these systems is that, they still don’t understand the complete language and complexity of the domain. You typically end up with custom domain adoptions and infinity language rules, which is definitely not smart enough to manage in longer run. The predictions of such systems are usually not accurate.

- Self learning chatbots - How often you have heard this terminology called self learning chatbots. This again is a misconception, where chatbots learns on its own. You have to train chatbot on what you what the chatbot to learn. Usually you would capture the user behavior details through their interaction with the chatbot application. This would include capturing user analytics information like capturing his likes or dislikes in some way, either through explicit or implicit means. Explicit information can be a user rating a product and implicit can be the time a user spent looking at a response.

Once you know the user well and have its data, its becomes a recommendation problem on what you want to recommend to the user. So, you end you building a recommendation algorithm to recommend something. For instance, for a fintech application, this would mean recommending similar stocks based on what stock he views regularly or his portfolio.

Different domain and use cases, would need different recommendation algorithms and that needs to be developed as part of the chatbot. However, the learning is boxed, for instance if you have a chatbot which assist you in booking restaurants, it can recommend you similar restaurants, but it can't recommend your places to stay, it only knows about your restaurants taste.

Well, someone can build a recommendation system which tracks what users eat and where they stay and then try to come up with a correlation which provide a recommendation, as the system now knows - "user eating XYZ, most likely are adventurous. So, recommend some trekking place." Again, in this case, recommendation is boxed on what you know and what you want to recommend. I don't know if any such hypotheses exist, but only through data and feedback that can be inferred. The point is, all of these hypotheses, data and feedback needs to be designed and developed, and saying chatbots learns on its own it's quite misleading.

- General purpose, generative chatbots – A chatbot which is capable of learning new concepts from scratch and provide responses like Human. As it learns from open domain, the chatbots would start behaving similar to the famous Microsoft Tay chatbot (which was forced to shut down on its launch day), as it started learning unwanted details from tweets and started posting inflammatory and offensive tweets. This is a classic example of what I quoted earlier - “AI can learn, but can’t think”. The generative chatbots are formulating the response based on probability of words and creating a grammatically correct sentence, without understanding the real meaning of it.

As I mentioned earlier, the first focus should be on getting domain specific chatbots right and with the current techniques we are far away from realizing the vision.

Will chatbots make human agents obsolete?

To answer this question, lets understand what functionality chatbots currently provide.

Current chatbot implementation do well for handling fixed set of dialogs with the user, repetitive tasks and certain initial aspects of customer service tasks. Wherever there is a fixed set of processes and flows to automate, chatbot can be used to provide 24 * 7 support for any queries. If human expertise is used for answering basic set of questions where answers are readily available, it would be eventually be replaced.

But in real-life scenarios, most of the conversation usually doesn't follow a fixed flow paradigm. But if the conversation moves from basic questions to questions which need further analysis, or the topic of conversation gets changed, you need a sophisticated chatbot implementation to take care of various conversation flows, identify the context switch, identify intents which your chatbot may not be aware of and create queries to find that information from your knowledge source. You are now moving from fixed set of flows to more dynamic flows which needs to be interpreted by your chatbot. Building such complex chatbot implementations requires sophisticated domain specific adoption using machine learning techniques and custom solutions. Current out of the box chatbot services fall short of building such chatbot implementations.

And even if you have all the data in the world at your disposal, infinite processing and computation power, using the current generation technology and research, you can never build a system that can compete with an expert human in the field. Taking even a 5 year horizon from now, I don't think we can develop such a level of intelligent chatbots.

For instance, can chatbots or an assistant, help doctor to recommend cancer treatments accurately and consistently. The answer is No.
The information provided from chatbot can aid doctors to take a clue from the answer provided, it may be right or wrong. You can never certify this. The chatbot would always act as assistance to an expert person to get some job done. Ultimately, these systems are throwing a bunch of answers based on some probabilities. The answers are limited to what you have fed into the system, you can't infer a new knowledge on the fly or can correlate information like a human expert to come to any conclusion.

While, there are research going on to use deep neural nets for conversation flows, we are still quite far away of building truly conversational interfaces which understands the nitty-gritty of language and domain. Also, the answers provided needed to be explainable and unless you have a way to backtrack on why a particular answer was provided, such deep neural systems can't be used for use cases which requires auditability and explainability.

In short, enjoy the smart chatbots that gives a perception on being intelligent, but intelligence is a long way away.

Can AI generate dynamic responses to questions

You can use deep learning to build a chatbot. Various deep learning architectures are available to solve specific variety of use cases. For instance, for computer vision (i.e. image recognition) you would use a convolutional neural network as the starting point, for language translation or text generation you would go with recurrent neural network and so on.

For understanding chat conversations, you would start with a variant of recurrent neural network. You will build a sequence to sequence model. A sequence to sequence model in simple words consist of 2 components, the first component (encoder) tries to understand
context of input sentence through its hidden layers and the second component (decoder) takes in the output from encoder and generates the response.

The above techniques require you to have a large set of training data, containing questions and responses. The technique works in a closed domain, but as the responses are dynamic in nature, putting it directly to your end users can be a bit risky. Secondly, these techniques don't work when you want to interpret the input sentence to extract the information and formulate a response on your own, like the shopping advisor query use case that we discussed above.

In case of open ended domain, the chatbots would start behaving similar to Microsoft Tay chatbot example I gave earlier.

Tip - With RNNs, the response/answer is dependent on its previous states (or earlier states). So, for deep conversational use case, where context needs to be available, the RNNs doesn't work. You need to employ variants on RNN called LSTM. (Long Short Term Memory networks). There are a lot of research going around this area. Going through various deep learning architecture is outside the scope of the book.

SUMMARY

The current generation of chatbots are weak form on AI, which offers an ability to understand the intent of the input message/question. In order for chatbot systems to understand the intent, it needs to be trained with the corresponding domain. You can ask the same question in multiple ways and the chatbot implementation can still infer the intent.
For dialogs, the current technology offers defining fixed conversation flows, so the interactions are boxed and finite.

Chatbots do well for managing productivity and certain aspects of customer service tasks. However, as the complexity of domain increases, current technology falls short, as even after sufficient training you would not get the required level of accuracy. You would need to rely on a combination of other machine language technologies and solutions like rules, inferences, custom domain metadata to get the solution delivered. These become a one-off solution, which becomes difficult to generalize. For some cases, even the one-off solution would be very complex, like building an advisor for recommending cancer treatments accurately and consistently.

While, there are research going on using deep neural nets, we are still quite far away from building a true conversational chatbot which understands the nitty-gritty of language and domain. Also, the answers provided needed to be explainable and unless you have a way to backtrack on why a particular answer was provided, such deep neural systems can't be used for use cases which requires auditability and explainability.

QUIZ - WHAT DO YOU THINK

To end the chapter, I would like to stimulate your brain by providing a set of questions, which you can hopefully answer with the information gained in this chapter

Quiz 1 -
Do you think a chatbot or an advisor currently exist that can assist doctors in providing correct cancer recommendations. Assume the system has been fed with every medical literature that exist in medical and cancer domain.

Quiz 2 -
Do you think a general purpose chatbot can be implemented with current technologies. Take example of Amazon, can the entire online experience for all departments be driven by a chatbot.

Quiz 3 -
When do you think chatbots would really be intelligent ?

Dear Readers, Thank you for Reading.

We hope you enjoyed reading the book, and the information provided would be a valuable resource in your AI journey.

I will be constantly be updating the book with chapters and latest information. This book is freely available on my website. For latest updates on the book, kindly visit -
http://navveenbalani.com/index.php/books/download-real-ai-a-definitive-handbook/

As an author, we strive for comments and feedback to improve our book and would greatly appreciate if you could leave your valuable feedback. Please share your review and feedback via email at
me@navveenbalani.com

Thank you.