

Case Studies of Machine Learning Applications in Retail, Hospitality, Education and Insurance Sectors

Mayank Chaturvedi, Khushboo Jain, Akshay Mehta, Tanmay Pal, Janhavi Sawal, Aishwarya Sharma, Sabina Tandon, Steffisara Varghese & Prasad Ramanathan¹
Dept. of Artificial Intelligence & Data Science; Mukesh Patel School of Technology, Management & Engineering NMIMS University, Mumbai, India

Abstract—Machine Learning & Deep Learning techniques have been put together in the form of “Lego Blocks” to create value-added solutions for customers as well as business owners across four different industry sectors, i.e. Retail, Hospitality, Education and Insurance. Chatbots and voice bots address multiple use cases like answering FAQ queries, providing guidance to users on specific transactions and invoking external web services. Pattern recognition & image recognition are shown to automate some tedious and error-prone activities. Topic modeling activities enable automated discovery of areas where new service offerings can be created based on customer inputs.

Keywords—Applications of Machine Learning, Deep Learning, Retail, Hospitality, Education, Insurance, Chatbot, Pattern Recognition, Image Recognition, Recommendation, Topic Modeling

I. INTRODUCTION

Machine Learning (ML) and Deep Learning (DL) techniques form the technology foundation for Artificial Intelligence (AI) applications like Predictive Analytics, Natural Language Processing (NLP), Pattern and Image Recognition. Commercial products such as Google DialogFlow, Microsoft Bot Framework, etc. use Supervised Machine Learning techniques, along with NLP, to enable chatbot implementations. Recognizing patterns in images is also maturing into a field of widespread commercial application. Maturity of unsupervised learning techniques has already been exploited by e-commerce companies that use such techniques in their recommendation engines.

By applying these techniques of supervised learning, unsupervised learning, natural language processing and pattern recognition, we can achieve significant gains in a variety of industrial sectors. In this paper, we present four such case studies where the very same set of AI/ML techniques are applied to accomplish substantially different use cases across four different industry sectors, i.e. Insurance, Education, Hospitality and Retail. Even though the use cases are different, the Information Processing Task that is inherent to each of these is substantially similar. As a business user, these technology enablers can be thought of as Lego Blocks, that can be put together to address a problem in that sector.

In the following sections, we detail out the use cases for which ML/DL techniques were applied in the four industry sectors (for both Supervised Learning as well as Unsupervised Learning situations). The business value of each use case is

demonstrated by implementing a prototype application for each sector. The benefits of such implementations for both end-users as well as the business owners are articulated. Details of the technology enablers are provided along with a discussion on how these implementations can be deployed in the real world.

II. RETAIL SECTOR APPLICATIONS

AI/ML technologies are shown to be of significant value in alleviating challenges faced by both the end-consumers using e-commerce platforms, as well as the administrators of e-commerce platforms through the following use case implementations.

A. Use Case Details

There are several e-commerce websites selling people items of their choice and necessity. More the sites, more the options to choose from, and more the challenge for the consumer to check through each site for availability, price, deals, recommendations, etc. A chatbot interface, that relies on NLP techniques, can significantly enhance the customer experience, and therefore, ensure customer-stickiness. Fig. 1 below shows how a chatbot can be integrated within the landing page of the e-commerce website.

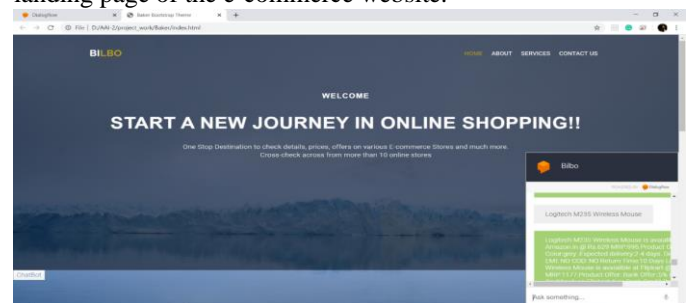


Fig. 1. E-commerce website with a chatbot integrated within its landing page

1) Responding to FAQs of customers

The most mature application of chatbots is to respond to FAQs from customers, such as store policies of purchase / return / refund, etc. From a business standpoint, this ensures that clients get responses to their queries at the “moment of truth”, when a transaction is likely to materialize. That way, consumers do not have to wait for responses from the help desk team of the e-commerce site. This also ensures that routine queries, which would have otherwise landed on the helpdesk phone support team, is responded in a Do-it-Yourself

manner. For e.g., a chatbot that is integrated with the store's inventory database can respond to queries on the availability of a product of a specific make/ model (e.g. availability of Amazon Alexa).

In addition to FAQ responses, chatbots can deliver greater value, when they are integrated with the enterprise landscape of the e-commerce company, and external web services via API calls. Some use cases where the chatbot can be used as the conversational User Interface to present results of such API calls are illustrated through the following implementations.

2) Price comparisons for customers

In Fig. 2 below, it is demonstrated how a user can query the chatbot for comparing prices of a product, such as mouse. On recognizing the intent of "Price Comparison", the chatbot invokes an external web service call through PriceYuge Data API [1]. This, in turn, obtains the price of this product from different e-commerce sites such as Amazon, Flipkart, etc. Details includes price, Product colour, Expected Delivery, COD availability, Product Offer etc

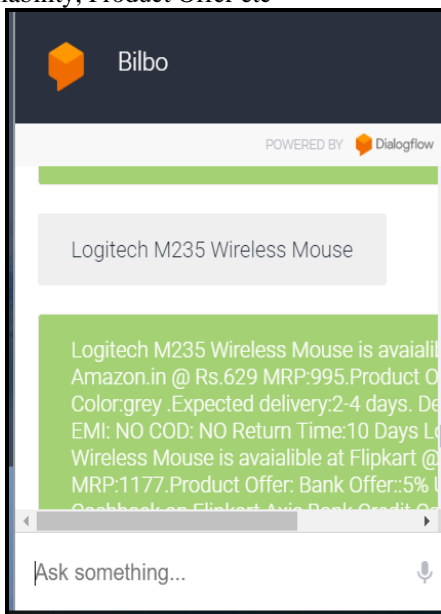


Fig. 2. Comparative Details of the product Logitech M235 Wireless Mouse fetched and shown in the chatbot window

3) Product Recommendations for Customers

Advanced statistical and machine learning techniques can be applied on some products to discover similarities across different products and visualize them in a more holistic manner. For e.g., in the context of cosmetic products, giving recommendations based on the chemical compositions of cosmetics [2]. Referring to the visualizations depicted in Figs. 3 and 4, the closer the distance between the two items is, the more similar the composition they have. Therefore, this enables consumers to compare the items without having any chemistry background.

4) Automated Product Labeling for Store Administrator

E-commerce websites need to ensure that products that come from different sellers on their platform are listed in a consistent manner. More the number of sellers on a platform, and greater is the number of items being listed, greater is the challenge of listing all features of the product correctly. Image recognition by applying Deep Learning algorithms can be

used to discover features of a product from a picture and ensuring that they get listed consistently. Thus, deep learning techniques can be used to auto-populate some of the common labels associated with a product. Fig. 5 below demonstrates the prototype implementation of how an image of a product can be uploaded, and Fig. 6 shows the various labels or features predicted for that product, which get automatically stored in the database. While the example below is that of a cosmetic product, similar results can be accomplished using images of clothing / accessories, etc. Thus, the process of populating the database is automated, efficient and less error prone. The system is trained on more than 3000 images and produces 97% accurate results.

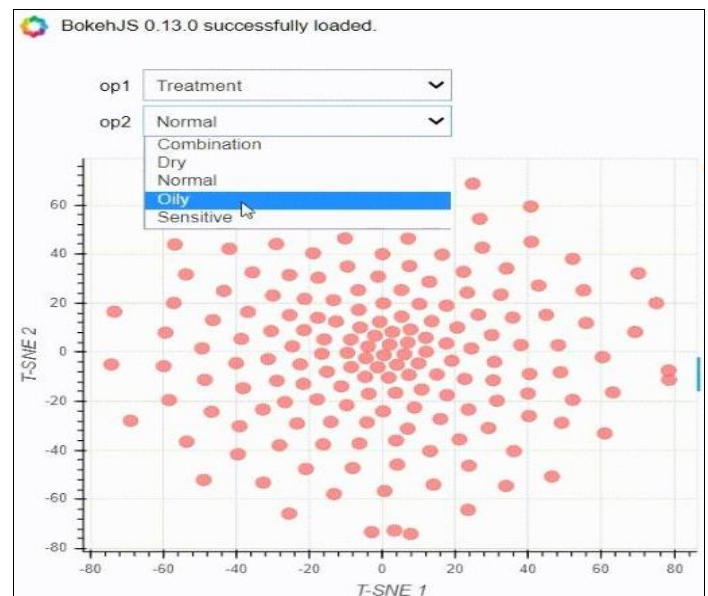


Fig. 3. Select the Skin Type from the options given.

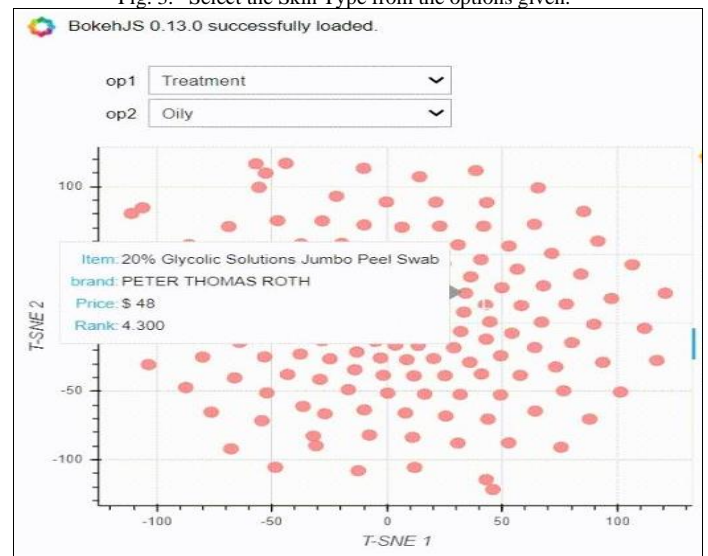


Fig. 4. Recommendation System recommends products with brand, pricing and Rank

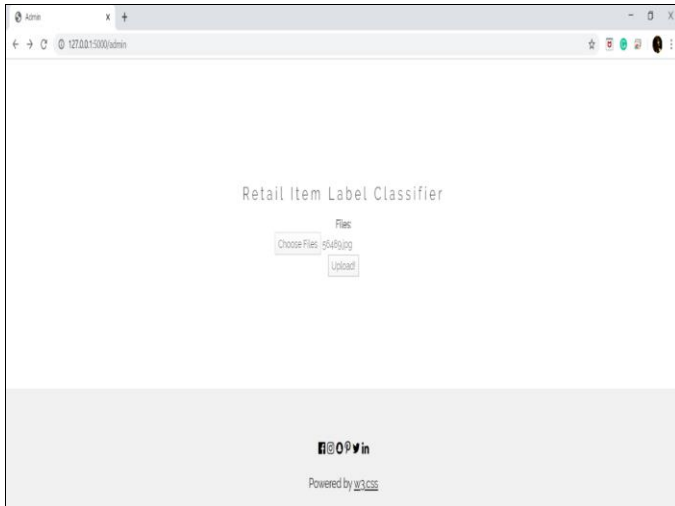


Fig. 5. Web support for Auto-labelling the product Image

5) Aid for updating FAQ knowledge base by Store Admin

Once a chatbot application is developed, it is imperative that the results of the interaction of the chatbot with end-users be analyzed frequently. By identifying those utterances for which the chatbot was unable to identify the intent, it is imperative that an administrator keeps the FAQs periodically updated and relevant to customer queries. Manually sifting through many queries that remained unanswered to discover those topics for which responses need to be configured in the chatbot can be a daunting task. Here, Topic Modelling on the unrecognized search queries to identify the clusters of topics that are frequently asked by customers, can enable clients to prioritize those areas where the chatbot needs to be updated with responses. The top topics thus discovered, can be included in future version of chatbot to serve users.

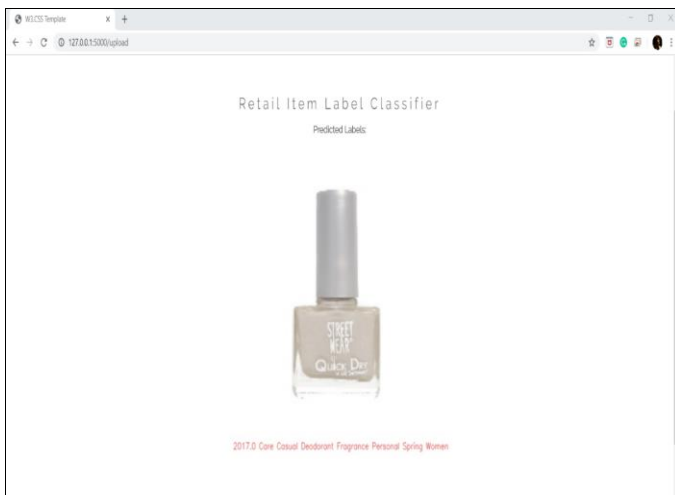


Fig. 6. Predicted labels from the input image are the red colored words


B. Technology Enablers & Implementation Details

The key technology enablers that were used in the implementation of this prototype application are listed below:

- Google Dialogflow [3]: The chatbot was built using Dialogflow which allows to create entities, intents, set contexts and training phrases for the product related FAQs and search queries.

- Ngrok: Web Tunneling tool used to provide local connectivity to the web based Chatbot Interface.
- Flask: A Python Framework used to provide webhook connectivity to the chatbot interface.
- MySQL Workbench: To provide database connectivity for storing users' feedback, unrecognized search queries, product and sales information.
- PriceYuge Data API: External API to make web service call to fetch product details from multiple e-commerce stores.
- NLTK Library: To perform text analysis on the user search queries.
- FastAI library [4], Product Image Dataset containing 40000+ images with label information for feature extraction from images
- Bokeh for Visualization library

Fig. 7 below shows the results obtained while training the Deep Learning model for identifying labels from images. Data for Deep Learning was obtained from Kaggle [5].



epoch	train_loss	valid_loss	accuracy_thresh	fbeta	time
0	0.409019	0.143534	0.965075	0.639964	00:43
1	0.216465	0.072474	0.972207	0.709925	00:39
2	0.136229	0.057026	0.975717	0.758519	00:39
3	0.097105	0.051267	0.977104	0.785743	00:39
4	0.075532	0.050152	0.977166	0.793437	00:39




Fig. 7. (a) Finding appropriate Learning rate for the multi-label classification model. (b) Training vs Validation loss plot obtained on training model for 5 epochs and on 3000+ images

III. HOSPITALITY SECTOR APPLICATIONS

In the context of the hospitality industry sector, AI/ML technologies can be applied to a variety of use cases of value to both end-users as well as website administrators, as detailed below.

A. Use Case Details

Suppose a person wants to plan a trip. For such a purpose, s/he would need details regarding suitable locations to go to, hotels to book for during their stay there, activities to plan during the trip, amongst other queries. To aid in such a planning process, a chatbot implementation could ideally be used as a Conversational User Interface to answer a plethora of customer queries.

1) Responding to FAQ type queries of customers

A chatbot was developed using Google Dialogflow to demonstrate how FAQ-type questions can be answered (Fig. 8), thereby relieving the helpdesk team to attend to greater value-add tasks. Such questions could include information about policies of booking/ cancellation/ refund, etc.

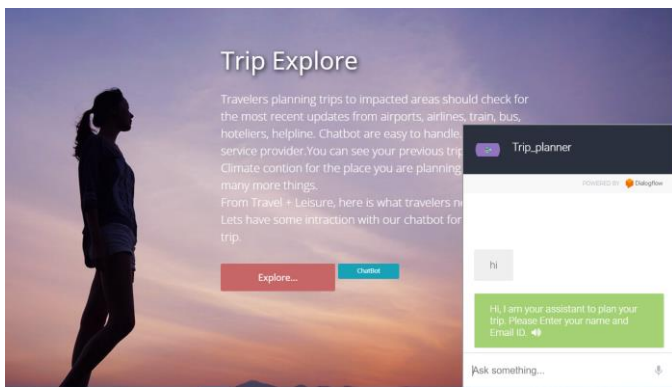


Fig. 8. Interactive chatbot as a pop-up within the web page

2) Getting destination weather information, map directions and setting calendar reminders for customer

By integrating the chatbot with external APIs such as weather, map, calendar it is possible for the user to get additional information about the trip that (s)he is planning to undertake by giving him/ her information about the weather at the destination, directions to the place and also adding a reminder in his/ her calendar for the planned trip. Thus, a chatbot's true value can be enhanced by such integrations with external web services, which help in personalized & contextual responses.

3) Getting information about prior trips for contextual responses to customers

By integrating the chatbot with the database of the tour operator, it is possible to figure out which locations the person has visited in the past. That way, the context of the user is used so that the same locations are not suggested to the person, when (s)he is interacting with the chatbot. Thus, interactions with the chatbot tend to be a lot more human-like in its responses, as its responses are contextual.

4) Auto-Labeling Type of Tourist Destinations for Tour Administrator

When a new hotel or property needs to be listed by the tour operator in their website, it is important that the facilities of the property are consistently stated. Getting to populate a

database with the right kind of information about a property (e.g. availability of swimming pool, gaming zone, gymnasium, etc.) can be a manual and error-prone task. However, pattern recognition from images using Deep Learning techniques can enable the tour operator to discover such features from images of the property and get the database populated in an automated manner. Fig. 9 below shows the results obtained by the implementation of such a Deep Learning algorithm to discover the features of a property. Thus, the task of the administrator gets significantly simplified.

5) Enabling speedy check-in

A QR Code can be generated for users who book a tour, so that their check-in at the destination property is fast-tracked. Such express check-ins into hotels and resorts, as well as for checking in into rooms, instead of using room key cards, can be an exceptionally great "delighter" for customers, while also reducing the routine tasks of the front-desk persons at a property. Fig. 10 illustrates the QR code generated for a client.

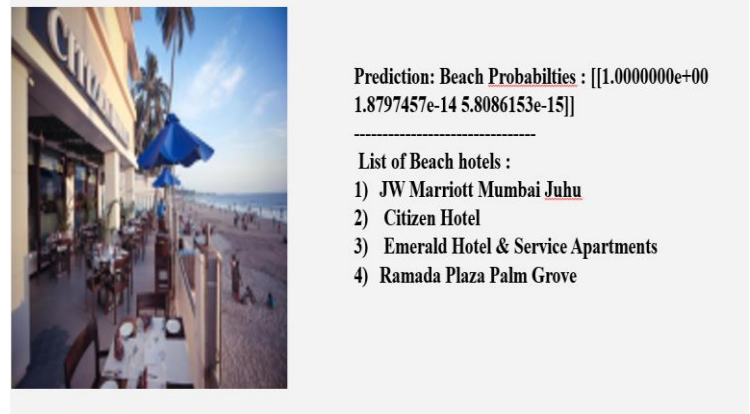


Fig. 9. Identifying property type as a "Beach Hotel" based on an image, and listing similar properties




Fig. 10. Providing QR Code to customers to enable speedy check-ins

6) Updating Offerings Based on Customer Interests

Identifying the queries for which the bot is unable to provide an answer to discover categories in them via Topic Modelling enables the Service Provider to identify domains or areas in which clients are expressing interest, for which the provider does not have any immediate offerings. This will help in growing the business of the hospitality service provider.

B. Technology Enablers & Implementation details

The chatbot was created with DialogFlow and embedded in an HTML website. Furthermore, Python was used for customizing interactions and MySQL to store and maintain

relevant databases. The complete interaction of the platforms was enabled using Flask as well as by Python libraries.

The chatbot was trained in DialogFlow to answer simple FAQs as well as interactive conversations. Advanced interactions such as user identification required the use of Python to extract user details from the chatbot conversations and to match it with the details stored in the MySQL databases via means of MySQL queries. The results generated after such a query were then customized in Python and displayed through the chatbot.

Python found primary application in customizations of responses, platform integration with MySQL, the Deep Learning algorithms used for Image Processing and Recognition, generation of QR codes and implementation of Topic Modelling for unanswered user queries.

External app interactions were also made available by the integration of Google Maps API, OpenWeather API [6] and Google Calendar API. All of these are open-source APIs and were thereby made use of for this project.

IV. EDUCATION SECTOR APPLICATIONS

In today's age, education can be provided to a wide variety of audience in all parts of the world using Internet-based delivery systems. In this case study, a few AI/ML technology enablers were used to demonstrate how they can enable students and administrators to get their work done faster and more efficiently.

A. Use case details

1) Chatbot for FAQ Responses

A chatbot based on NLP techniques was developed to provide responses to FAQs raised by both prospective students as well as previously registered students. For e.g., prospective students can get answers pertaining to courses in various fields and of various durations. Similarly, previously registered students can get answers about any course-related information in the chatbot. Such chatbots can be readily extended to respond to voice-based queries from students, which is a far more natural way of interaction, than text-based interaction.

2) Setting up reminders in calendar

While setting reminders in a calendar for upcoming events like webinars, exams, etc. can be done in a stand-alone manner, using the conversational user interface of the chatbot to get this done provides a single interface from which this can be done. This requires integration of a chatbot with the Google Calendar APIs.

3) Topic modeling to discover topics mentioned in a student's answer

In order to enable automated grading of subjective answers provided by a student in an answer paper, it is possible to use Topic Modeling to identify clusters of topics that the student's answer includes. This will enable faster and more error-free grading of answer papers submitted by students. Similarly, queries raised by students to the chatbot can be grouped together into topics.

4) Facial Recognition of Students

Students in front of their computers can send their images using their webcam to the website of the e-learning portal. Using Deep Learning algorithms, it was demonstrated how facial recognition can be applied to discover if the image

belongs to a previously registered student or a new student. This can be useful in situations where it is necessary to figure out who is present in front of the computer. If there is a requirement that a student must mandatorily attend some part of a webinar/ e-learning course, then the web cam can be used to send the photograph of the person sitting in front of the computer. Facial recognition techniques (using Deep Learning algorithms) can then be used to discover if the student who is expected to be in front of the computer is truly present or not. Fig. 11 below shows the results of the facial recognition features implemented in this project. The confidence factor in identifying the images can be improved with greater amount of training data.

B. Technology Enablers & Implementation Details

The key technology enablers used in this implementation are as follows:

- Google Dialogflow – used to develop the Chatbot Interface.
- SQLite3 was used to store data of - student, courses, topics etc.
- Flask web application framework was used for Fulfillment of Dialogflow using Python.
- ngrok was used to create secure tunnels between Dialogflow webhook and localhost
- Selenium web driver was installed to open the website and look for the respective search.
- All webhook fulfillment code was development using Python.
- Facial Recognition for identifying the Student(user) using Resnet50 architecture of Keras library.



Fig. 11. Facial Recognition Output

V. INSURANCE SECTOR APPLICATIONS

In the insurance sector, several transactions right from requesting for a quote to filing insurance claims are getting digitized. With the technologies of AI/ ML reaching a higher level of maturity, some of the use cases where insurance sector can benefit from these technologies are highlighted in the following paragraphs.

A. Use Case Details

1) Chatbot for FAQ responses

Chatbot provides a conversational UI that enables end-users to interact with the insurance company's portal using Natural Language. Such a chatbot can readily answer some FAQ type queries related to the different insurance products offered by the company. Additionally, contextual queries pertaining to a specific customer can be answered when the chatbot is integrated with the insurance company's database to let the user proactively know when his insurance payment is

due, or the status of an insurance claim that was filed by the user. An example of such an interaction of the user with the chatbot is shown in Fig. 12 below.

2) *Image recognition to identify damaged automobile parts*

During the process of filing an insurance claim, end-user is often concerned about the approximate value of the damaged part, and if it will be covered by the insurance policy that (s)he has for the automobile. In this context, having image recognition techniques to discover the type of damage to the vehicle can help in estimating the cost of repair/ replacement. As shown in Fig. 13 below, a picture of the damaged part of the vehicle can be submitted to a Deep Learning model. The model can identify if the picture has “features” like a dent, or a shattered windshield or that of a damaged headlight. By integrating this feature within the insurance portal, an end-user can readily discover the cost of repair/ replacement and decide for himself/ herself, whether it is worth filing an insurance claim and risk losing the “No-Claim-Bonus” discount.

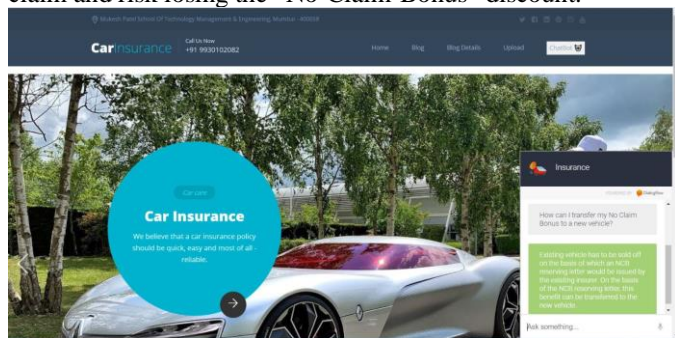


Fig. 12. Chatbot responding to customer’s queries

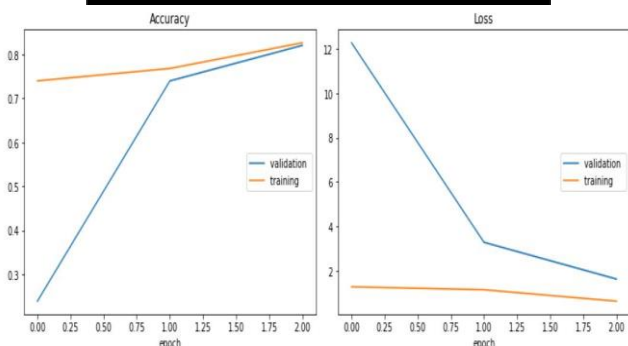


Fig. 13. Predicting a dent on an automobile based on Image Recognition

3) *Verifying details in an insurance claim*

After an insurance claim has been filed by the customer, if the client has indicated that the accident was caused during a rainy day, it is readily possible for the insurance company to check if that was true. For this, an invocation of the Weather

API, from within the portal can provide on-line answers. Specifically, this can be further eased if the interaction with the chatbot invokes the weather APIs to provide answers to the insurance company on the true weather condition at a given place at some time in the past. This will help in eliminating any fraudulent claims too.

B. *Technology Enablers & Implementation details*

The key technology enablers that were used in the implementation of the above use cases include:

- MySQL Database to store structured information about the customer
- Checking the current weather status using “openweathermap” API
- Keras ResNet50 for image recognition, along with FastAI
- Topic modelling on unrecognized intents to get the 3 topmost queries of the customers.
- Google’s Dialogflow – Used to define several Entities, Context and Intents for the Chatbot.
- Ngrok - To create a secure tunnel on your local machine along with a public URL
- Flask – A python framework used to integrate Internal/ external web services calls with the Chatbot

VI. RESULTS & CONCLUSIONS

Through the implementations of different use cases across different industry sectors, it has been demonstrated that chatbots / voice bots can be exploited in multiple ways. A summary of the ways in which chatbots can be beneficial across industry sectors is provided in Table I below.

TABLE I. CHATBOT APPLICATIONS				
Type of Use of Chatbot	Retail Sector	Hospitality Sector	Education Sector	Insurance Sector
<i>Answering FAQ type queries (generic responses)</i>	Information about store timings, policies of refund/ return	Information about tour packages/ hotels	Information about courses, fees, pre-requisites	Information about different types of insurance coverage
<i>Enabling Transactions by integrating with enterprise applications</i>	Guiding the user to place an order	Guiding the user to make a reservation	Guiding the user to register for an on-line course	Providing an insurance quote; Guiding the user to file an insurance claim
<i>Value-added services by invoking external web service calls</i>	Scraping data from other retailers’ websites to do price comparisons; setting calendar reminders about sale dates	Invoke weather API to inform user about current weather conditions at a destination; setting calendar reminders about upcoming trip	Set calendar reminders for date / time when the course is scheduled to begin; or when fee payment is due	Set calendar reminders about premium payment due dates; get weather conditions at the time when an accident is reported

Unsupervised learning technique of Topic Modeling was applied on the text of queries raised by customers to a chatbot, which the chatbot was unable to provide any satisfactory responses. This gives the business owners an opportunity to discover those service offerings that are under-served.

Pattern recognition and image recognition have been demonstrated to be of significant business value across all the sectors as summarized in the Table II below.

TABLE II. PATTERN & IMAGE RECOGNITION APPLICATIONS

<i>Retail Sector</i>	<i>Hospitality Sector</i>	<i>Education Sector</i>	<i>Insurance Sector</i>
(a) Automated discovery of features of a product based on an image for listing in e-commerce platform (b) OCR the label of a product to discover ingredients, and recommend the right one, e.g. right cosmetic for right skin tone	(a) Automated discovery of features of a property to be listed by an aggregator based on photos provided by a property owner, e.g. swimming pool, gym, etc. (b) Suggest properties with certain features based on user preferences (c) Generate QR code to enable speedy identification and check-ins at hotels	(a) Discover whether the person trying to register for a new course is a new student or an existing student (b) Identify if persons attending on-line courses are truly the ones sitting in front of the computer	(a) Discover the make and model of an automobile to provide an insurance quote (b) Based on a picture of the damaged vehicle, identify the possible damage and provide an estimate of repair cost

By integrating all three of the above technology enablers from within a website, significantly high value-added services are shown to be provided, enabling customers to have a richer experience and simultaneously reducing the efforts from business owners.

REFERENCES

- [1] PriceYuge, <https://price-api.datayuge.com>
- [2] Jiwon Jeong, For Your Skin Beauty: Mapping Cosmetic Items with Bokeh, <https://towardsdatascience.com/for-your-skin-beauty-mapping-cosmetic-items-with-bokeh-af7523ca68e5>, Nov. 28, 2018 [Accessed in April 2020]
- [3] DialogFlow, <https://dialogflow.com/>
- [4] FastAI, <https://www.fast.ai/>
- [5] Kaggle Data Set, <https://www.kaggle.com/bhaskar2443053/fashion-small>
- [6] OpenWeatherMap, <https://openweathermap.org/>