



Machine Learning-based Natural Language Processing for Turkish Venue Recommendation Chatbot Application

Gorkem Toprak¹, Jawad Rasheed^{2*}

¹ Department of Computer Engineering, Istanbul Aydın University, Istanbul 34295, Turkey, gorkemtoprak99@gmail.com

^{2*} Department of Computer Engineering, Nişantaşı University, Istanbul 34398, Turkey, (ORCID: 0000-0003-3761-1641), jawad.rasheed@nisantasi.edu.tr

(İlk Geliş Tarihi 17 Mayıs 2022 ve Kabul Tarihi 28 Ağustos 2022)

(DOI: 10.31590/ejosat.1117635)

ATIF/REFERENCE: Toprak, G. & Rasheed, J. (2022). Machine Learning-based Natural Language Processing for Turkish Venue Recommendation Chatbot Application. *Avrupa Bilim ve Teknoloji Dergisi*, (38), 501-506.

Abstract

In recent years, mobile applications occupy an important place in our lives. Thanks to machine learning and the chatbots advancements, people can easily access the information they want or the things they need. Our purpose in this study is that users can access detailed information about the main venues of Istanbul from a single application, access their menus and photos, read the blog posts written by gourmets about these restaurants, show nearby venues according to their instant location, and to gather young gourmets chasing different tastes under a single application by developing a Turkish chatbot that can find the most suitable venues according to users' wishes or offer suggestions about venues to visit. In order to evaluate our application, we tested it on both Android and iOS platforms and achieved successful results on two platforms.

Keywords: Foods, Venues, Places, Smartphones, Android, iOS, Chatbot, Flavor Hunter.

Türkçe Mekan Öneri Chatbot Uygulaması için Makine Öğrenimi Tabanlı Doğal Dil İşleme

Öz

Son yıllarda mobil uygulamalar hayatımızda önemli bir yer tutmaktadır. Mobil uygulamalar ve içerdikleri chatbotlar sayesinde kişiler istedikleri bilgilere veya ihtiyaç duydukları şeylere kolaylıkla ulaşabilmektedir. Bu çalışmadaki amacımız, kullanıcıların tek bir uygulamadan İstanbul'un belli başlı mekanları hakkında detaylı bilgiye ulaşabilmeleri, menülerine ve fotoğraflarına ulaşabilmeleri, gurmelerin bu restoranlar hakkında yazdıkları blog yazılarını okuyabilmeleri, yakındaki mekanları anlık konumlarına göre gösterebilmeleri ve kullanıcıların isteklerine göre en uygun mekanları bulabilen veya gidilecek mekan önerilerini sunabilen bir Türkçe chatbot geliştirerek farklı tatlar peşinde koşan genç gurmeleri tek bir uygulama altında toplamak. Uygulamamızı değerlendirmek için hem Android hem de iOS platformlarında test ettik ve iki platformda başarılı sonuçlar elde ettik.

Anahtar Kelimeler: Yemekler, Mekanlar, Akıllı Telefonlar, Android, iOS, Chatbot, Lezzet Avcıları.

* Corresponding Author: jawad.rasheed@nisantasi.edu.tr

1. Introduction

Today, with the effect of the pandemic, people's desire to eat tasty but hygienic food and drink in quality places has reached an important dimension. The pandemic has greatly affected all food service businesses and they have closed their shops for a long time. This mobile application has been developed to bring together young gourmets who have adopted the food and beverage culture and to help places even if it is small.

While the number of smartphone users has exploded, in 2021, technology research firms invested \$1.5 trillion in digital transformation technologies, with estimates that this figure will rise to \$1.8 trillion by 2022 (Samuels, 2022). In addition to this, innovative apps may now be applied in a variety of ways thanks to smartphone technology. Researchers developed various applications such as smart attendance (Rasheed et al., 2019) and blood donation system (Fahim et al., 2016) to provide hustle-free management. In the past, people had to go to more than one website to get information about restaurants and they spent a long time. However, thanks to Lezzet Avcilari, people will be able to get detailed information about the restaurants in the Istanbul region via their mobile phones and see the restaurants close to them on the map.

The Lezzet Avcilari application is a social media platform, but an application whose main subject is restaurants-spaces. In this application, people will be able to get detailed information about restaurants in Istanbul, see the atmosphere, menus, and comments made by other users, and save the current restaurant, and if they want, they will be able to read blog posts written by gourmets about that restaurant. Users will be able to register and log in to the application in a very simple way. If the user has entered his email and password correctly, the application will redirect him to the main screen. On the home screen, users will be able to access the trending restaurants according to their weekly scores, access the categories and reach the restaurants in that category, and finally, read the featured articles written by gourmets. They will be able to access all the restaurants in the application from the venues page and click to reach the page with detailed information about the restaurant. They will be able to access the restaurants nearby with the map on my nearby screen and, if they wish, they will be able to examine the region in more detail by switching between the places on the map. On the blog screen, they will be able to read the articles written by gourmets about restaurants and if they liked the article, they will be able to press the like button and leave feedback. On the profile page, users will be able to access the restaurants he has returned to, edit their profile, see the restaurants he has saved and access the settings of the application. In addition, users will be able to message each other through the application, so that young gourmets will be able to recommend quality places to each other and stay in constant communication. If one of the users who are messaging is not active in the application, they will be able to receive notifications with the message sent by the sender. And they will be able to access this notification they receive by accessing the notifications page on the main page. They will also be able to receive notifications of trending restaurants or new blog posts thanks to firebase cloud messaging within the application. Lastly, users will be able to get information about restaurants thanks to the bot on the home page, so they will not need to switch to more than one page, more precisely, users will be able to communicate with the robot

via a mobile application and determine the region and venue features in line with their preferences.

Most artificial intelligence systems have long-term aims of understanding humans and engaging with them to make their lives simpler (Wu et al., 2020). These are services that operate in a specific region and provide skilled technical solutions for certain business sectors. Recommendation systems that focus on people today provide the most logical choices to users by evaluating data from many sources in a variety of industries, including e-commerce, education, and health. Artificial intelligence-based systems speed up the finding of the best techniques and contents, especially when enormous volumes of data are accessible (Wen et al., 2015).

This mobile application, it is aimed to develop a Turkish chatbot that provides information and suggestions about food places in Istanbul. The task-oriented chatbot (Akma et al., 2018) is designed to provide both automatic and interactive written responses to user queries in Turkish using natural language processing and deep learning methods (Elifoglu & Gungor, 2018). By communicating with the chatbot through a mobile application, users can determine the region and venue features according to their preferences. It will provide a more user-friendly service to consumers via chatbot (due to conversation functionality). The places to be visited would satisfy consumers' expectations due to the exact choices that may be specified ahead of time (Keeble et al., 2020). Customers and venues benefit economically and socially from these types of apps. Instead of spending time researching to find a dining venue, people would save this time to do something else. Furthermore, clients from various places would be able to communicate globally. Although the application will keep some personal data (such as location, socio-economic status, etc.), we don't anticipate any direct legal impact.

When we examine the use case of the study, people seeking eating places can benefit from a more helpful system that integrates with booking, travel, and food-ordering sites (Kim et al., 2020). It may be used as a stand-alone application to quickly investigate meal options. On the other hand, When we examined the success criteria of the study, we thought that we could obtain the following items:

- A dialogue system that helps users feel as though they're conversing with real individuals.
- Generating human-like phrases
- Users are guided appropriately and accurately by the chatbot system.
- Retrieval of a collection of meal options based on the preferences of the users.

The rest of the paper is divided into four sections. Section 2 outlines the relevant applications. Section 3 describes the suggested design and implementation details along with test results. Finally, Section 4 concludes the paper.

2. Materials and Methods

In this section, we are going to discuss the methods and frameworks we used in this study.

2.1. Dialogflow

Dialogflow is Google's natural language understanding tool, which allows your chatbot to comprehend language. Dialogflow

tries to comprehend the user's intent based on the training sentences you provide. Dialogflow knows how to automatically react to inquiries from your users based on those training phrases. It was first released as an API. All of the capabilities you'd expect are included in Dialogflow, including:

- The ability to extract user intents from each given speech using natural language processing (NLP).
- Machine learning is used to educate your bot to provide increasingly useful replies.
- Scalability is achieved by using Google services that are close by. Google Cloud Platform is where Dialogflow is hosted.
- Integration with a variety of chat applications and messaging systems.

2.1.1. Intents

During a communication, an intent categorizes a user's intent. Several purposes can be assigned to each agent. Dialogflow intents will match the end-user phrase with the best intent in your tool when the user speaks something called an end-user expression. Goal categorization is another name for goal matching. An example of the intents in our study is given in Figure 1.

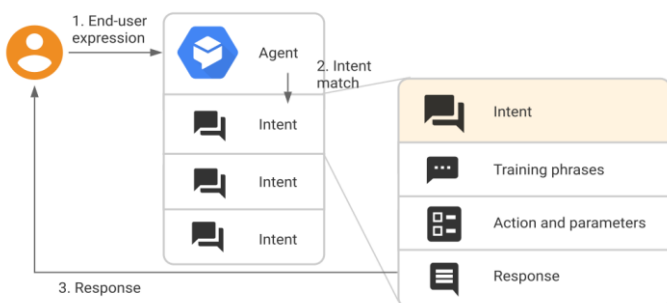


Figure 1. Simple dialogflow intent's structure

2.1.2. Training Phrases

End-user expressions, often known as training phrases, are hypothetical phrases that end-users could write or utter. We may phrase the same thing in a variety of ways, but it all refers to the same duty. For instance, if we want to find a restaurant to dine at, we may say, "List me the nearby restaurants or, if I'm hungry, please advise a place where I can eat". These are known as training expressions, and they are used to fit the user's wishes, as shown in Figure 2.

☞ günaydin
☞ hey
☞ nasilsin
☞ selam naber
☞ merhabalar
☞ selamlar
☞ selam
☞ Merhaba

Figure 2. Training phrases of the project

2.1.3. Responses

We need to get back to the user when we've completed all of the processes. We use responses to provide this feedback, as shown in Figure 3. In the dialog flow console, we may set up several replies. Platform-specific responses can also be configured.

Text Response	
1	Selam! Size nasıl yardımcı olabilirim?
2	Merhaba! Size nasıl yardımcı olabilirim?
3	Günaydın, size nasıl yardımcı olabilirim?
4	Enter a text response variant

Figure 3. Responses of the study

2.2. Seq2Seq Learning Method

While researching the proposed approaches, we found a method that can be useful to us. The name of this method we found was the Seq2Seq learning method. The seq2seq learning approach is the answer for sentence creation. Two components are employed in the seq2seq model: encoder and decoder. Figure 4 depicts the architecture of seq2seq. The encoder receives the user's input sequence and passes it to the decoder. It seeks to forecast the next step for each step based on the sequences generated up until that point. The decoder creates an output sequence based on the data it receives. The input is not restricted to the length of the output in the solution procedure. The lengths of the input and output sequences might be variable. In addition, if the system's input sequence is too lengthy, the system may fail. To avoid this, RNN structures are utilized, and each step of the encoder sequence is saved in a vector. This vector will be transmitted to the decoder when the encoder section is completed.

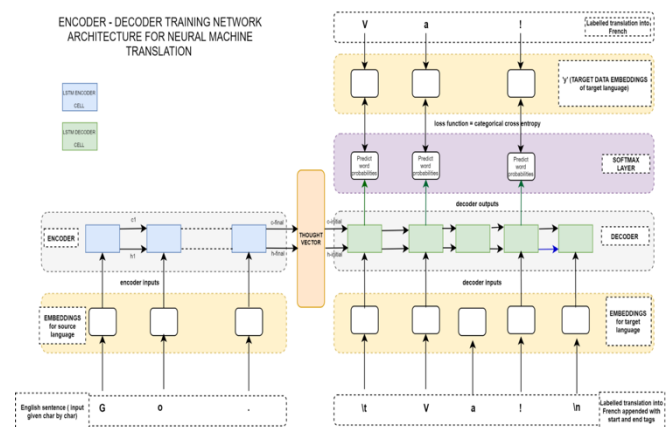


Figure 4. Architecture of Seq2Seq method

2.3. Flutter

The interactive mobile user interface is implemented via Flutter. Flutter is a suite of user interface tools from Google that lets you build engaging, embedded apps for mobile, website, desktop, and embedded devices with a single codebase. Flutter is a cutting-edge framework you can feel! It's a lot easier to construct mobile apps with it. The documentation for Flutter has a lot to offer, and it's all fairly extensive with easy examples for basic use cases. To the question of why we should use flutter, it is explained in short as follows:

- You can publish on 2 platforms with a single code system, that is, it provides development experience on a single platform for iOS and Android.

- Good documentation is essential for getting started with new technology. You can learn a lot from Flutter's tutorial library and it's all very insightful and meaningful with easy examples for basic use cases.
- By using Flutter's system, you can get useful and nice-looking structures.
- Thanks to the Hot-Reload feature, you can see the changes you have made to the code on your application in a short time.

More detailed information about it can be found in the next section.

3. Results and Discussion

The results we have obtained as a result of the work we have carried out are listed separately below.

3.1. Tests/Results of the Chatbot

The performance assessment models generated phrases that were identical, comparable, or wholly different from the content expressions provided in the venue descriptions. Figure 5 shows some examples of correct and erroneous sentences generated by the natural language generation component.

<p>Correct Sentence:</p> <p>Dialog Act: (type = 'inform', name: 'Pasto', near: 'Venedik Pastanesi', customer_satisfaction: 'Yüksek')</p> <ul style="list-style-type: none"> • Reference Sentence 1: Venedik Pastanesi'nin yakınlarında bulunan Pasto , müşterilerin memnun kaldığı yerlerden biridir . • Reference Sentence 2: Pasto , yüksek müşteri memnuniyetine sahip olup Venedik Pastanesi yakınında hizmet vermektedir . • Generated Sentence: Venedik Pastanesi yakınında bulunan Pasto , müşterilerinin çok memnun olduğu bir mekandır. <p>Incorrect Sentence-1:</p> <p>Dialog Act: (type = 'inform', name: 'Hiç Hane', near: 'Macarons d'Antoinette', region: 'Karatay, Konya')</p> <ul style="list-style-type: none"> • Reference Sentence 1: Hiç Hane , Konya Kartay'da Macarons d'Antoinette restoranına komşudur . • Reference Sentence 2: Hiç Hane , Konya Karatay'da , Macarons d'Antoinette'in yakınındadır . • Generated Sentence: Bursa Çukurova'da , İskender Lokantası yakınında Çiya Sofrası bulunmaktadır . <p>Incorrect Sentence-2:</p> <p>Dialog Act: (type = 'inform_not', customer_satisfaction='Yüksek', cuisine = 'Deniz Mahsulleri ve Dünya Mutfağı Yemekleri', price_range = 'Ucuz', region = 'Anadolu Hisarı Merkez, İstanbul')</p> <ul style="list-style-type: none"> • Reference Sentence 1: İstanbul'da Anadolu Hisarı Merkez'de , deniz mahsulleri ve dünya mutfağı yemekleri sunan yüksek müşteri memnuniyetinde bir mekan yoktur . • Reference Sentence 2: İstanbul'da Anadolu Hisarı'nın merkezinde müşterilerini memnun bırakan , Deniz Mahsulleri ve Dünya Mutfağı Yemekleri sunan bir mekan bulunmaktadır . • Model Output: İstanbul Anadoluhisarı Merkez bölgesinde bulunan deniz mahsulleri ve alkolü içecekler ile alkolü içecekler sunan bir mekan yoktur .

Figure 5. Correct and incorrect sentences

Furthermore, the sentences with errors were generated by the model. These mistakes can be divided into two categories. The structure and content of the first wrong statement are accurate, however, the field values in the sentence are incorrect in the conversation act. There are structural and contextual issues with the second wrong sentence. Delexicalization is done on the input dialogue acts before forwarding to the natural language generation component to overcome these issues (Li et al., 2018). Figure 6 shows the natural language understanding component's projected purpose and entities for a text. The component forecasts the top intent, as well as its entities and other potential intentions, and examples of discussion flow with user statements and chatbot behaviors.

Figure 7 shows examples of discussion flows with user statements and chatbot behaviors. Also, below are examples of both good and bad conversations. As the bad dialogue demonstrates, the chatbot can reply to unconnected sentences. We are going to work on the training set and the technology to increase the chatbot's accuracy.



Figure 6. Proposed application dialog workflow

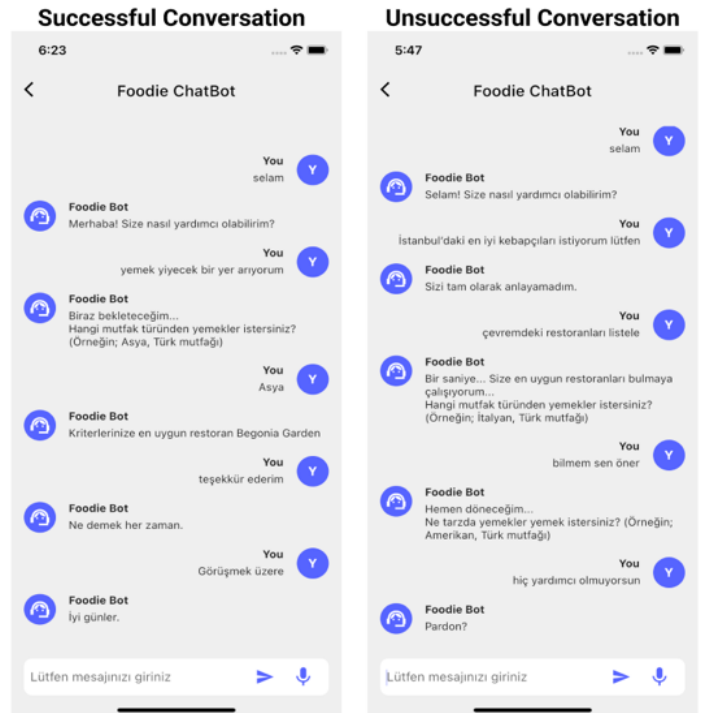


Figure 7. Successful and unsuccessful conversation

3.2. Tests/Results of the Mobile Application

When designing a mobile app project, tests are conducted to verify or validate whether it meets user needs. For these and other reasons, mobile app testing is critical. It is critical to test mobile apps to ensure that they function properly. Manual or automated testing of the application on mobile devices is possible. The lower the quality of a mobile application means that the less fixed code is written on the device, the better it is. In other words, if the dependency is fully dependent on your service layer and managed from here according to the changes, the application will be quite manageable. An authentication module, a restaurant review module, creating a new restaurant feedback module, a map module, review blogs, and chat are all included in the test plan.

flow of the application and have been faithfully advanced in design.

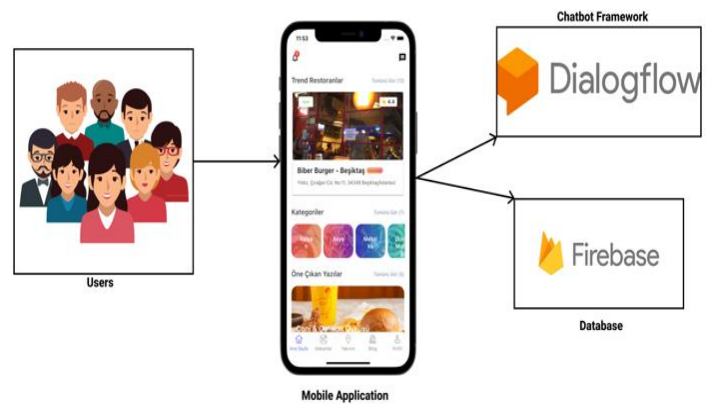


Figure 8. Successful and unsuccessful conversation

The mobile application is very simple to use. Users do not need any technical information; all they need to do is to get information about the restaurants they are interested in, whether through a smart chatbot or by navigating through the application. Many users can log in and use the application simultaneously. For users to communicate with the smart chatbot, it is expected that they will first send a message. Because this bot, which has been trained, prepares its answers according to the messages written by users. If the chatbot does not understand the message written by the user, in this case, it responds by entering the fallback intent and telling the user that the fallbacks it has received are like "I don't understand, please repeat". And lastly, the implementation is highly flexible and fault-tolerant, also scalable in the future, and easy to maintain and add new features.

3.3.2. Chatbot

Chatbots are a sort of computer software that mimics and processes human speech (Akma et al., 2018), allowing users to converse with digital devices as if they were conversing with actual people. Chatbots can range from simple programs that can always respond to a simple question with a one-line response to digital assistants that learn and grow as they collect and analyze data. Chatbots follow a set of restrictions that limit their ability to reply to specific predetermined words or orders that have no meaning. As a result, if you don't utilize the suitable term while interacting with the bot, the chatbot may get stumped and react with "sorry, I don't understand, please repeat" responses. You may converse with chatbots as if they were actual people since they are powered by artificial intelligence. Because artificial intelligence allows them to comprehend the language and reply appropriately. Furthermore, chatbots get smarter over time as they learn something from each interaction with users. Besides, the Dialogflow framework is used to create the chatbot. To train the chatbot, a dataset was created from human-annotated Turkish conversational data in the restaurant reservation domain (Oflazlar & Saraçlar, 2018). Task-oriented (declarative) chatbots are one-purpose systems that specialize in a specific task. They create replies to user inquiries using rules, NLP, and a small amount of machine learning. Interactions with these chatbots are particularly well-structured. Task-oriented chatbots may answer typical queries in everyday life, such as questions regarding basic procedures. These are the most popular chatbots right now. As there are no Turkish dialogs available to train a dialog management component, we first examine the E2E dialog dataset of Turkish dialogs in the domain of the venue (Sharma et

- The authentication module will be the focus of the first test. The goal of the test is to ensure that none of the authentication processes fail. Users must double-check their email addresses and passwords.
- The second test will primarily focus on the process of users receiving restaurant reviews. The goal of the test is to guarantee that the user receives a unique result. Users will see a restaurant list that includes information such as the restaurant's name, photos, menus, and operating hours
- The third test will mainly test the module for generating feedback by users on the restaurant's detail screen.
- The fourth test will test whether users can view restaurants within the application via a Google Maps module. Users will be able to see a restaurant list that includes information such as the restaurant's name, photos, rates, and phone numbers.
- The fifth test will essentially test whether users can access blogs written by gourmets.
- The last test will test whether users can message privately, whether the message is read or not, and whether a notification is sent by the message user.

3.3. Discussion on Working of Proposed Approach

After researching the issues in the related works and techniques mentioned in the previous section, we recommended the architecture of the mobile application in Figure 8. We have divided the architecture into 3 main parts, taking into account the hardware requirements, the demands of the application development platform for smartphones, interaction, and task-oriented chatbots. These 3 main parts are; mobile application, task-oriented chatbot, and database. oriented chatbots.

3.3.1. Mobile Application

The system interacts with users via a mobile application (Vinaik et al., 2019). The mobile application is implemented in Flutter/Dart. Users interact with our stylish and design-friendly mobile application that displays a list of restaurants to access the details of their dining venues and for reliable information. Users are required to register when they first download the application. Registration and login processes are quite simple. Those who have successfully logged in can start using our modern and aesthetic application, try restaurants and dishes of different cultures, communicate with a task-oriented chatbot, read blog posts and message different users. Venues that match the user's choices are gathered and shown on a map based on the preferences obtained through the mobile app. Users may traverse the map and set their search criteria, as well as read one-sentence descriptions of locations (Isabela et al., 2018). Moreover, all of the screens are carefully linked with each other to understand the

al., 2021). We compile and implement training data for our chatbot by examining these provided dialogues and extracting dialogues other than the purpose of our application.

3.3.3. Database

The application's database is built on Firebase. Firebase is a NoSQL cloud storage solution for collections and documents. Because the application must adapt to a variety of use cases, it is well suited for usage as the application database. This is not the case with a structural database like a SQL database. Cloud Firestore is essentially a cloud-hosted NoSQL database for storing and synchronizing data. It, like Firebase Realtime Database, uses real-time listeners to keep the data in sync across client apps and provides offline support for mobile and web, allowing you to create responsive apps that run regardless of network latency or Internet access. Firestore is used to store all the data of the user account, restaurants, blogs, comments, and chats. Firebase Cloud Storage is utilized to store larger files such as photos and videos. Cloud storage is utilized in the mobile application to save all of the photographs of users, and restaurants to store images. Moreover, it is preferred for security reasons provided by Google when uploading or downloading files from our Firebase-connected applications.

4. Conclusion

This study highlights our attempts to create a task-oriented conversation system in Turkish for venue recommendation. The Dialogflow framework is used to create natural language comprehension and conversation management components. People nowadays have a difficult time deciding where to eat depending on their unique tastes. The main purpose of this study is to provide a user-friendly, human-like agent that interacts with people to solve the time-consuming problems people face and find places to eat. It will be an innovation for both individuals and restaurants to develop an application that tackles this problem. Also, another purpose of the study is to provide a little help to places that have been closed for a long time due to the pandemic. Lastly, The mobile application has been developed and tested on actual devices and within the internet context, it functions without a hitch. This is only a beta version, and there will be more updates in the future.

References

- Akma, N., Hafiz, M., Zainal, A., Fairuz, M., & Adnan, Z. (2018). Review of Chatbots Design Techniques. *International Journal of Computer Applications*, 181(8), 7–10. <https://doi.org/10.5120/ijca2018917606>
- Elifoglu, M., & Gungor, T. (2018). A restaurant recommendation system for Turkish based on user conversations. *2018 26th Signal Processing and Communications Applications Conference (SIU)*, 1–4. <https://doi.org/10.1109/SIU.2018.8404153>
- Fahim, M., Cebe, H. I., Rasheed, J., & Kiani, F. (2016). mHealth: Blood donation application using android smartphone. *2016 Sixth International Conference on Digital Information and Communication Technology and Its Applications (DICTAP)*, 35–38. <https://doi.org/10.1109/DICTAP.2016.7543997>
- Isabela, E., Drona, J., Fadhilah, N., Tanoto, D. F., Harefa, J., Prajena, G., Chowanda, A., & Alexander. (2018). NYAM: An e-ISSN: 2148-2683
- Android Based Application for Food Finding Using GPS. *Procedia Computer Science*, 135, 393–399. <https://doi.org/10.1016/j.procs.2018.08.189>
- Keeble, M., Adams, J., Sacks, G., Vanderlee, L., White, C. M., Hammond, D., & Burgoine, T. (2020). Use of Online Food Delivery Services to Order Food Prepared Away-From-Home and Associated Sociodemographic Characteristics: A Cross-Sectional, Multi-Country Analysis. *International Journal of Environmental Research and Public Health*, 17(14), 5190. <https://doi.org/10.3390/ijerph17145190>
- Kim, H., Jung, S., & Ryu, G. (2020). A Study on the Restaurant Recommendation Service App Based on AI Chatbot Using Personalization Information. *International Journal of Advanced Culture Technology*, 8(4), 263–270. <https://doi.org/10.17703/IJACT.2020.8.4.263>
- Li, X., Wang, Y., Sun, S., Panda, S., Liu, J., & Gao, J. (2018). *Microsoft Dialogue Challenge: Building End-to-End Task-Completion Dialogue Systems*. <http://arxiv.org/abs/1807.11125>
- Oflazer, K., & Saraçlar, M. (Eds.). (2018). *Turkish Natural Language Processing*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-90165-7>
- Rasheed, J., Alimovski, E., & Rasheed, A. (2019). mManagement: Wi-Fi Hotspot based Attendance Application using Android Smartphone. *2019 1st International Informatics and Software Engineering Conference (UBMYK)*, 1–5. <https://doi.org/10.1109/UBMYK48245.2019.8965588>
- Samuels, M. (2022). *What is digital transformation? Everything you need to know about how technology is changing business*. Zdnet. <https://www.zdnet.com/article/what-is-digital-transformation-everything-you-need-to-know-about-how-technology-is-reshaping/>
- Sharma, R., Dhir, A., Talwar, S., & Kaur, P. (2021). Over-ordering and food waste: The use of food delivery apps during a pandemic. *International Journal of Hospitality Management*, 96, 102977. <https://doi.org/10.1016/j.ijhm.2021.102977>
- Vinaik, A., Goel, R., Sahai, S., & Garg, V. (2019). The Study of Interest of Consumers In Mobile Food Ordering Apps. *International Journal of Recent Technology and Engineering*, 8(1), 3424–3429. <https://www.ijrte.org/wp-content/uploads/papers/v8i1/A9219058119.pdf>
- Wen, T.-H., Gasic, M., Mrkšić, N., Su, P.-H., Vandyke, D., & Young, S. (2015). Semantically Conditioned LSTM-based Natural Language Generation for Spoken Dialogue Systems. *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 1711–1721. <https://doi.org/10.18653/v1/D15-1199>
- Wu, C.-S., Hoi, S. C. H., Socher, R., & Xiong, C. (2020). TOD-BERT: Pre-trained Natural Language Understanding for Task-Oriented Dialogue. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 917–929. <https://doi.org/10.18653/v1/2020.emnlp-main.66>