



# Virtual Assistants for the Cultural Heritage Domain

Aleš Tavčar<sup>(✉)</sup>, Jernej Zupančič, and Matjaž Gams

Jozef Stefan Institute, Ljubljana, Slovenia  
ales.tavcar@ijs.si

**Abstract.** Virtual assistants and similar software tools are gaining importance among phone and computer users. The most well-known assistants (Siri, Cortana, Google, etc.) provide general information to users and cannot be adapted to specific needs. Custom implementations usually cover specific domains and are specialized to provide a comprehensive set of information or functionalities from that area. Currently, there is a lack of general applications for cultural heritage, since most implementations are specialized virtual guides for museums or exhibitions. We present an overview of virtual assistants that we developed and are applied to the area of cultural and natural heritage, where we try to improve the current shortage of tools. The presented prototypes provide various functionalities to users and can help them to discover, learn and plan visits to cultural sights.

**Keywords:** Virtual assistants · Natural language processing · Web services

## 1 Introductions

In recent years, virtual assistants (VA) have become a popular tool that helps users find information faster and in a more accurate manner. They are web-based services that serve as a human-computer interface to deliver information from a specific domain that the user is interested in or to interact with an information system. Companies like Google, Microsoft or Apple developed general assistants for mobile devices that help users find restaurants or shops close to their location, manage their phone or provide computer-aided support for software packages.

The main advantages of virtual assistants are their ability to understand, to a certain extent, natural language and present the answers in an enriched format (additional links, multimedia presentation, etc.). This way, users can provide questions and queries in natural language, which makes virtual assistants well-suited for less skilled or handicapped users. In addition, virtual assistants can implement speech-to-text and/or text-to-speech systems to further improve the user experience.

Nowadays, assistants are being developed for a wide variety of domains and applications and they thrive due to the increasing publicly available data on the internet, both structured and in free-form. In addition, web based services that offer API access can be integrated into the assistants to provide a comprehensive user experience (for example, booking of flights, transportations, events, tours, medical visits, hair-dressers, etc.). In addition, modern third party services that provide advanced

functionalities that can also be managed through APIs (for example, messaging, text mining, recommendations, path planning, etc.) can also be exploited by assistants.

Applications of virtual assistants in the cultural heritage domain are not as widespread as one would expect. There is a notable lack of applications for the interested public, providing general information and covering non-specific domains, since most applications are custom solutions intended as virtual guides for museums and exhibitions. The introduction of more generic applications and virtual assistants would help popularize and spread the appreciation for cultural heritage also among youngsters.

In this paper we provide an overview of the various virtual assistants developed at our institutions for the domain of cultural heritage and sightseeing. Virtual assistants integrate several advanced functionalities. Text mining and Natural Language Understanding (NLU) to identify the intent of the users and extract important entities from text. Next, they implement API calls to third party services in order to provide trip planning functionalities and data fusion. All presented assistants are available online and can be modified and thus applied to new domains or different languages.

## 2 Related Work

Virtual assistant in the domain of cultural heritage have been already applied to virtual visits of cultural sites. They are deployed in real or virtual environments as virtual guides to engage the visitors and deliver a comprehensive learning experience [1]. Other implementations also provide mechanisms, such as multimedia, storytelling, and enriched content to entertain the user during the learning process [5, 6]. All these agents must be flexible, responsive, and provide natural interaction in order to be seen as a social partner by visitors.

Several applications were developed that deliver a personalized museum guide. Mathias et al. [2] proposed a new method for personalized museum tour recommendations. Their research tackles the problem of optimizing museum visits according to visitor's preferences and artwork importance. Huang et al. [3] developed a museum guide system that uses association rule mining to discover recommendations from both collective and individual visiting behaviors. Pechenizkiy and Calders [4] presented a simple user-focused framework for personalizing museum tours that is focused on efficient learning since the system should be able to quickly provide relevant suggestions only after a small set of user preferences.

Other general applications of virtual assistants are in the area of education for instructions and learning [8], delivering information from large data sets like encyclopedias in libraries [9], for informational purposes in e-commerce [10], merging heterogeneous data using multi-agent conversational agents [11], tourism [12], and several others.


### 3 Overview of Deployed Virtual Assistants

#### 3.1 Platform for VA Creation

We developed a cloud-based platform Asistent [2, 7] for the creation of virtual assistants and during a national project we created and integrated several VAs for web pages of municipalities in Slovenia. The developed web platform supports the creation, management and integration of virtual assistants to an arbitrary web page. It was developed as a Software as a Service (SaaS), where clients access the service via a web browser, which eliminates the need for software and hardware maintenance. Since the software runs in the cloud all future upgrades are applied centrally and are available to all clients instantaneously.

The knowledge base of the deployed VAs for the municipalities was constructed manually and contains the information that municipalities deemed the most important. It covers the general information regarding the municipality, their acts and regulations, instructions and news for residents, natural and cultural heritage and touristic information.

Figure 1 shows an example of a VA developed for the Jozef Stefan Institute. The VA runs on the web pages of the institute and delivers related content to its visitors. In addition, the VA is enriched with specific applications that can be accessed through the Applications bar (see Fig. 1). They are an extension or addition of the basic functionality, since they offer the use of external services. For example, ticket booking for museums and payment of various fees. Moreover, applications can be defined as categories that set the context and focus the search of the assistant to some specific content. For example, users might want to search for answers regarding squares or buildings only in the tourism category. Each provider can define his own set of applications that are displayed in the virtual assistant.



**Fig. 1.** One instance of the VA with the application tab enabled.

### 3.2 Virtual Museum Guide

Museums offer to its visitors a wide selection of artwork at their premises. The vast amount of exhibitions is often overwhelming for the visitors and most of them might not be of interest to them. This might cause to miss artwork or entire exhibitions that are best suited to their personal preferences. The growth of the Semantic Web in recent years and the availability of more and more content in the digital and linked data format available on the internet supports the creation and reuse of services of a higher added value. One of such is the virtual assistant for museums that offers virtual tours of institutions covered in the Google Arts & Culture project (Google Street View). At the time of writing this paper, more than 3400 museums and sights are included in the Google Arts & Culture project. To deliver a simple proof-of-concept we mapped into the virtual assistant content from three museums (British museum, Art Institute of Chicago, and National Gallery of Australia).

The application provides a natural language interface using a virtual assistant created through the Asistent platform. The users can input questions or requests and the VA takes care of identifying the correct answer and displays the corresponding feed from Google Street View. In addition, a simple recommender system is integrated in the web service that provides suggestions regarding exhibits targeted for a specific user. Since the recommender does not get implicit rating for items from users, we had to design a recommendation system that learns user preferences based on implicit data (queries, time spent on each answer, motion, etc.) and item features. The algorithm implements a content-based filtering mechanism that is well-suited for this kind of domains. To generate a recommendation, the recommender does not use data from other users but the similarity between items in the dataset and the history of previously searched items to predict ratings for user-item pairs not present in the dataset. We showed that in the limited domain, with a small number of learning examples and user-item pairs the recommender works well, even with new users.

The virtual assistant can deliver information about specific exhibit rooms in the museum (for example, Egypt room), groups of similar objects (for example, Chinese weapons), or important exhibits (for example, the Rosetta stone). The virtual assistant contains the most relevant items and groups from the three museums, but it does not provide a comprehensive overview of the museum, since the mapping and creation of the knowledge base was performed manually, which is relatively time consuming.

The virtual assistant displays a description in the text area (bottom part of the assistant window in Fig. 2) and shows the related feed from Google Street View in the background page. The user can move around at his discretion and explore the presented exhibit. The user can also select the context in which the answer should be presented. For example, he might want to see the Egypt pottery exhibition in the British Museum, not from the other two mapped museums. In this case, the user selects the appropriate museum from the Application dropdown (the top left part of the virtual assistant in Fig. 2). This gives the user the ability to focus the search and obtain more relevant results.




Fig. 2. A query about the Rosetta stone in the British Museum.

### 3.3 Slovenian Cultural Heritage Database

The developed virtual assistant was created with the aim to deliver a wide range of information to users regarding cultural, natural heritage, and visiting sites in Slovenia. Due to the large amount of data and information that needs to be provided, the mechanism for content delivery was designed in such way to remove the need for manually inputting and editing database entries. This is a crucial step to enable a comprehensive delivery of content for all sights in a specific country. We implemented a web crawler that parsed several web pages related to cultural heritage and touristic information in Slovenia (for example, [www.slovenia.info](http://www.slovenia.info)). The crawler was able to classify the type of information that a page has and the type of the attraction (castle, lake, museum, etc.), extract metadata (location, region, descriptions, photos, related web page), and store the obtained information in a structured way. In addition, the crawler tried to obtain additional content also from external links that were included in the provided web pages. All the extracted content was stored in a relational database and the text in the main fields were indexed using Whoosh [3], a python library for indexing free-form or structured text and quickly find matching results based on simple or complex search criteria. Since the text in the database fields is in free-form, Whoosh provides an ideal solution to search within the text and at the same time obtain a relevance score that can be used to sort the search results.


The entire search mechanism in the virtual assistant is composed of four parts. First, the NLU module is used to obtain the intent of the user. The intent can be regarded as

the type of the question. The user might want to search for a specific sight or obtain a list or related items. Second, entities are extracted from the query. They are linked to the information extracted with the web crawler. For example, the type of the sight (museum, castle, lake, churches, etc.) and the region or location where the sight is. Next, the obtained values for entities are matched against the records in the database and a set of results is obtained based on the search criteria. Finally, the answer to be displayed in the virtual assistant is generated. A short description is provided if the user searched for a specific sight, or a list of matched items if the query return multiple hits.

The VA prototype supports the following types of questions by a user:

- Show me some lakes in the Dolenjska region.
- List the castles near Kranj.
- Show me some info regarding the Ljubljana castle.
- Museums in Ljubljana.

All the answers are provided in text form, with additional formatting when providing the description of a sight – a short description in the text are below the virtual assistant (see Fig. 3), an url link to the detailed description, category and a location of a sight, and the related web page displayed in the background.



**Fig. 3.** The virtual assistant provides information regarding the Predjama castle.

In Fig. 3 the user searched for a specific sight (the Predjama castle). The assistant identified the adequate entity and obtained the stored entry in the indexed database, based on the title field of the sight. A short description is provided in the text area of the virtual assistant and the related web page is displayed in the background.

The results of a different query are presented in Fig. 4. In this case, the user searched for lakes in the Gorenjska region of Slovenia. The NLU component identified two entities in the query. First, it classified the type of the query as a list of items. Next, the type of the sight as lakes and the region as Gorenjska. The virtual assistant then searched



for all items in the database that matched the identified entities (the region field as Gorenjska and sight type as lakes). Since the query type is a list of items, the module returns the first ten results with the highest relevance score. The list is presented in the text area of the assistant and each element is a hyperlink to the related sight that is shown in the same text area, with the corresponding web page displayed in the background.

The developed assistant and the search mechanism can be applied to different domain and languages.



**Fig. 4.** The virtual assistant showing a list of lakes in the Gorenjska region.

### 3.4 Virtual Assistant for Sightseeing

Tourism is a domain that can support the introduction of virtual assistants and they can provide significant added value to tourists. Companies that organize trips, tourist centers in cities and others can use the functionalities of a virtual assistant as additional services that they offer. Virtual assistants can deliver information from an area or regarding specific sights, provide content in a segmented, appealing, and entertaining format, they can help users to book tickets for events, sights, and urban transportation, create personalized city tours, and more.

We are currently developing a comprehensive platform that provides all the listed functionalities to enable smart tourism. The system uses several external services to improve the tourist's user experience during trip planning and provides guidance amid the actual trip. The platform integrates: (i) Rocket.Chat, a chat platform where users can communicate; (ii) Interface with Messenger; (iii) eTulist [4], a tour planner based on a recommendation system; (iv) Asistent, a rule based question-answering toolkit; (v) Dialogflow, a natural language understanding toolkit.

The prototype version of the platform supports various kinds of interactions and queries during the communication with the user. The classification of user intents is performed by the specialized virtual assistant that can understand queries and commands from the following categories: general information about a sight, sights near a specific

place, similar sights near a specific place or region, adding sights to the preference list, generate path to a sight, recommendation of sights. In addition, the virtual assistant is able to recognize an important word or sequence of words in the query. These words are from the following categories: sight name, sight type, place, and region.

Based on the recognized inputs from the user, the NLU module constructs the appropriate commands, methods, and API calls. The webhook functionality then retrieves the needed data from third party services, if required.

In Fig. 5 part of the conversation between a user and the virtual assistant is presented. The user asks the assistant to recommend something from the Gorenjska region. The search mechanism of the assistant works similarly to the one presented in Subsect. 3.3. The intent of the user is identified and the appropriate entities are extracted from the query. In the presented example, the intent is to get a recommendation and the entity region is Gorenjska. The assistant responds with two lists, the first is a list of potentially interesting places and the second is a list of sights that might be of interest to the user. The user can select one of the presented options by clicking on the interactive buttons or write a new query. The system provides a short description of the selected sight and again lists all possible actions linked to the sight. As before, the user can click

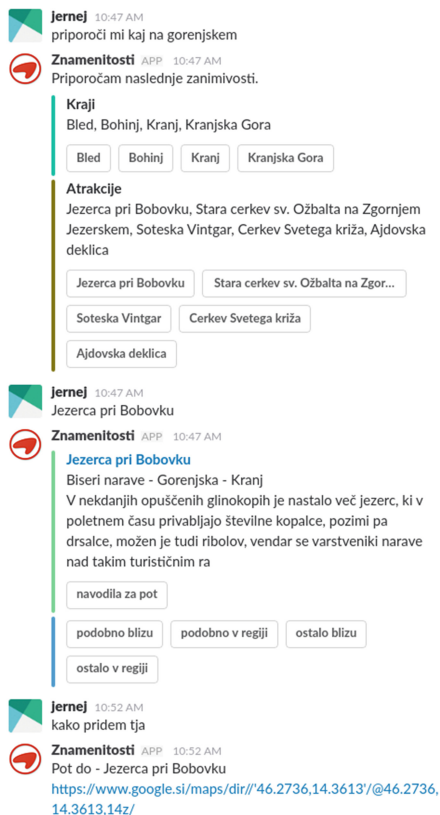
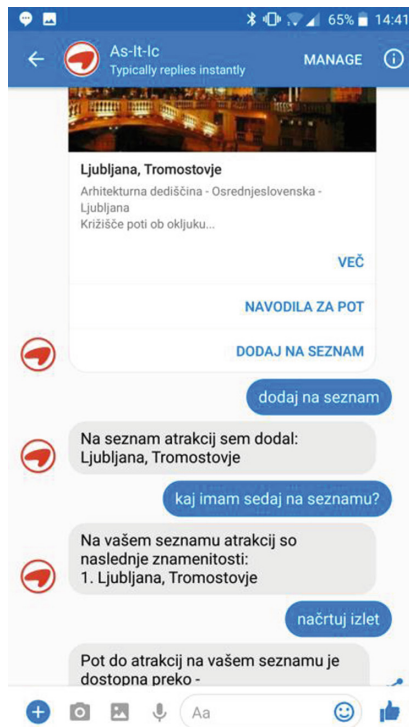


Fig. 5. Part of the communication flow with the virtual assistant.



on one of the buttons or write a new query. In the provided example, the path to the sight is computed and presented in the form of a link to Google Maps.

Figure 6 presents an example of a conversation in the Messenger interface. The user searched for the Tromostovje sight (a bridge in Ljubljana). He then wants the system to add the sight to his bookmarks. Next, the user wants the system to list all his bookmarks. Finally, he instructs the system to plan a trip based on the sights present in the list. The system collects the items and sends a request to the external service eTurist (tour planner) through the API. eTurist computes the shortest path between all items, considering all limitations (transportation method, time, lunch, budget, etc.), recommends and includes new sight along the way that might be of interest to the user, and send back the response in a JSON object that is used to construct an answer by the virtual assistant. Part of the response is also a graphical representation of the computed path in Google Maps, with marked locations and a short description for each sight.



**Fig. 6.** The virtual assistant interface in the Messenger client.

## 4 Conclusions

This paper presents an overview of the various virtual assistants developed at our institutions for the domain of cultural heritage and sightseeing. We presented a service for the creation of virtual assistants, an advanced tool for virtual interfaces and three prototypes of virtual assistants providing functionalities and content for different domains. First, a web based virtual guide for museums that delivers information about specific items, groups of exhibits, and rooms inside the British museum, Art Institute of Chicago, and National Gallery of Australia. The assistant shows a short description for each query and displays the pertinent feed inside the museums from Google Street View. Second, a virtual assistant for delivering information about sights, cultural and natural heritage in Slovenia. The assistant implements an intelligent mechanism for the identification of intent and extraction of entities from the user query. This way it is possible to deliver specific and correct information to the user. Finally, we presented a comprehensive platform for smart tourism and sightseeing. The platform integrates chat functionalities in several interfaces, a NLU module, a virtual assistant to understand the requests of the user and deliver enriched answers, and is able to use several third party services to obtain specific information and create personalized trips.

All services and prototypes are available online and can be applied to other, different domains and languages.

## References

1. Bickmore, T., Pfeifer, L., Schulman, D.: Relational agents improve engagement and learning in science museum visitors. In: Vilhjálmsson, H.H., Kopp, S., Marsella, S., Thórisson, K.R. (eds.) IVA 2011. LNCS (LNAI), vol. 6895, pp. 55–67. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-23974-8\\_7](https://doi.org/10.1007/978-3-642-23974-8_7)
2. Mathias, M., Moussa, A., Zhou, F., Torres-Moreno, et al.: Optimisation using natural language processing: personalized tour recommendation for museums. In: Proceedings of the 2014 Federated Conference on Computer Science and Information Systems, pp. 439–446 (2014)
3. Huang, M.Y., Liu, C.H., Lee, C.Y., Huang, Y.M.: Designing a personalized guide recommendation system to mitigate information overload in museum learning. *Educ. Technol. Soc.* **15**(4), 150–166 (2011)
4. Pechenizkiy, M., Calders, T.: A framework for guiding the museum tours personalization. In: Proceedings UM 2007 International Workshop on Personalization Enhanced Access to Cultural Heritage (CHIP) (2007)
5. Cimiano, P., Kopp, S.: Accessing the web of data through embodied virtual characters. *Semant. Web* **1**(1), 83–88 (2010)
6. Doswell, J.T.: PECA: pedagogical embodied conversational agents in mixed reality learning environments. In: AIED, p. 957 (2005)
7. Kuznar, D., Tavcar, A., Zupancic, J., Duguleana, M.: Virtual assistant platform. *Informatica* **40**(3), 285–290 (2016)
8. Graesser, A.C., et al.: AutoTutor: a tutor with dialogue in natural language. *Behav. Res. Methods* **36**(2), 180–192 (2004)

9. Rubin, V.L., Chen, Y., Thorimbert, L.M.: Artificially intelligent conversational agents in libraries. *Libr. Hi Tech* **28**(4), 496–522 (2010)
10. Xu, B., Pan, Z.G., Yang, H.W.: Agent-based model for intelligent shopping assistant and its application. In: *The first Conference on Affective Computing and Intelligent Interaction*, Beijing, Citeseer, pp. 306–311 (2003)
11. Eisman, E.M., Navarro, M., Castro, J.L.: A multi-agent conversational system with heterogeneous data sources access. *Expert Syst. Appl.* **53**, 172–191 (2016)
12. Čereković, Aleksandra, et al.: Implementing a multi-user tour guide system with an embodied conversational agent. In: Liu, J., Wu, J., Yao, Y., Nishida, T. (eds.) *AMT 2009*. LNCS, vol. 5820, pp. 7–18. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-04875-3\\_7](https://doi.org/10.1007/978-3-642-04875-3_7)
13. Project Asistent - virtual assistant for municipalities. <http://www.projekt-asistent.si/>. Accessed 26 Apr 2018
14. The Whoosh Library. <http://whoosh.readthedocs.io/en/latest/intro.html>. Accessed 26 Apr 2018
15. Cvetkovic, B., et al.: e-turist: an intelligent personalised trip guide. *Informatica* **40**(4), 447 (2016)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

