Questions and Answers: Important Steps to Let AI Chatbots Answer Questions in the Museum

Stefan Schaffer[1(✉)], Aaron Ruß[1], Mino Lee Sasse[1], Louise Schubotz[2], and Oliver Gustke[2]

[1] German Research Center for Artificial Intelligence (DFKI), Berlin, Germany
stefan.schaffer@dfki.de
[2] Linon Medien, Berlin, Germany
chim@linon.de

Abstract. In this paper, we describe our work within the research project "CHIM - Chatbot in the Museum". CHIM is an AI-based chatbot prototype that enables conversational interaction using text and speech input: visitors can ask questions about certain artworks and receive answers in multimodal formats (text, audio, image, video). The application will be tested in the Städel Museum, Frankfurt/Main, Germany. To develop a proper Natural Language Understanding module, we adapted an existing categorization approach, gathered visitor questions, and structured them into twelve distinct content types. The preliminary results suggest that our approach to subdivide the previously overloaded content type *meaning* into further categories was successful, leading to a more balanced distribution of the data. We further describe the Natural Language Processing mechanisms employed here; these follow a multi-tiered approach using techniques like Rasa, BERT, and cosine-similarity to generate answers with different degrees of effort. Future steps are the implementation of dialog management, the refinement of the NLP strategies by integrating additional answers for selected exhibits, and the implementation of the final layout and interaction design. We are planning to test and evaluate the CHIM prototype on site in the Städel Museum in late 2021.

Keywords: Chatbot · Conversational interaction · Digital Museum Guide

1 Introduction

In recent years, artificial intelligence, or AI, has gained increasing attention in museums all around the world. In this context, AI was initially mainly used

We want to thank the Städel Museum Frankfurt for their support. This research is part of the CHIM project of the research initiative "KMU-innovativ: Mensch-Technik-Interaktion", which is funded by the Federal Ministry of Education and Research (BMBF) of the Federal Republic of Germany under funding number 16SV8331.

M. Wölfel et al. (Eds.): ArtsIT 2021, LNICST 422, pp. 346–358, 2022.
https://doi.org/10.1007/978-3-030-95531-1_24

for example for image recognition or database analysis in general, which can be considered as "classical AI domains" [8]. Today, we find a wider range of applications in museums using AI, utilizing neural networks, machine learning, robotics, computer vision, deep learning, or natural language processing. Moreover, there are exhibitions focusing on AI itself.

Our approach in the research project "CHIM - Chatbot in the Museum" is to use AI within a chatbot museum guide application aimed at improving visitors' museum experience. The chatbot AI should be able to answer specific questions about certain artworks and thereby help to eliminate known pain points in knowledge transfer and learning situations in museums: Often, a personal human guide is not available at a given time. Additionally, in the current pandemic situation, tour groups that cluster in front of an artwork are no longer allowed. Digital media guides, which allow for a more individualized experience, normally offer only "one-way-information" such as audio guide texts but cannot reply to specific questions. According to the concept of "free-choice learning" [10], the kind of learning that occurs in museums fundamentally differs from the type of learning that happens in schools. Whereas in schools, one is forced to learn content that is not self-selected, in the museum, you can choose to learn about objects and artworks that interest you. This kind of learning is described in a "contextual learning model" [10]. One key factor to improve free choice learning is to keep visitors activated and personally involved with the story. To accomplish this, content delivery has to take into account visitors' motivations, expectations, and their personal leisure values by giving them a maximum of choice and control in how they want to learn about an artwork. Our chatbot allows visitors to ask any questions they have about an artwork. On the one hand, this open question functionality complicates finding an appropriate answer by the AI. On the other hand, it allows us to offer information tailored to the specific users' interest at that particular moment. The chatbot could become a sort of "virtual guide" that is available at any time or place to answer visitors' questions. In contrast to a guided tour, where visitors could shy away from asking "stupid" questions in front of other group members, these social context barriers are usually lower when interacting with a machine. Compared to traditional media guides, a chatbot allows visitors to ask the questions that interest them at that moment. They do not have to choose from predefined content. In this way, we hope to simplify the learning process, boost visitors' attention, and ultimately increase visitors' satisfaction. In addition, by evaluating visitors' questions and interactions with the chatbot, museums will be able to improve their educational offers, since they will learn more about what visitors want to know. To make the CHIM chatbot available to a wider audience, we intend to implement it as a smartphone application. Unlike some previous chatbot applications developed for museums [11], the system is not specifically aimed at attracting younger audiences, but ideally caters to museum visitors of diverse ages and backgrounds.

We developed the current version of CHIM to be used in the Städel Museum, Frankfurt/Main. This has two main reasons. Firstly, the museum has recognized the importance and promoted the use of innovative digital applications in the cultural heritage sector for several years now and has set up a team specifically dedicated to digital aspects of their educational agenda. We are extremely grateful to the museum and its staff for their kind support in the development of CHIM and the hosting of the on-site evaluation of the prototype in late 2021. Secondly, we have access to a large corpus of audio guide texts, written and produced by Linon Medien specifically for the artworks exhibited at the Städel Museum. As we will elaborate in the following sections, in the CHIM project we explore whether we can use these existing texts in order to find answers to visitors' questions.

Regarding previous work, we want to point out theoretical approaches, especially in the field of digital humanities. Some scholars postulate that digitalization and the massive application of AI technologies could lead to new methods in analysis and rating patterns in art history [13]. A content-focused chatbot AI allows us to gain insights into topics such as user-generated content. Furthermore, these insights can provide important impulses for the discussion about the sovereignty over the interpretation of art and cultural heritage.

With respect to relevant technical aspects, well-known chatbot and dialog platforms like Alexa (Amazon), Dialogflow (Google) and others need to be mentioned: they enable intention detection for many fields but are not sufficiently "case sensitive". If one asked Google questions of the kind that we collected and evaluated (see Sect. 3), regarding a specific artwork, one would get internet and Wikipedia hits, but not necessarily a proper answer. However, the number of AI-based conversational guiding systems, specialized in the field of cultural heritage or museums is growing [5]. A wide variety of approaches can be found, starting from systems that provide audio or media guide information via platforms like WhatsApp by typing numbers [2], to more conversational chatbot applications [1].

The goal of CHIM is to develop a learning, multimodal dialog system for knowledge transfer in museums. While working towards the envisioned chatbot, we explored different methods for making the system understand visitor questions and for finding suitable answers, e.g., by extracting the answers from existing audio guide texts. In this paper, we describe the steps we undertook in building a Natural Language Understanding (NLU) model for the classification of visitor questions. Adopting an approach from [6], we identified distinct content types for questions asked about selected artworks from the Städel Museum and developed Natural Language Processing (NLP) strategies for generating answers by using these content types, complemented by additional annotations. One novel contribution of CHIM to the field is that the system allows for user generated questions, rather than relying on pre-scripted dialogues, as other German language museum chatbots currently do [3,4]. Moreover, the advantage of developing our own NLU und NLP models as opposed to relying on for example Dialogflow, is that it enables us to store and process our data in accordance with German data protection laws, a non-trivial aspect of the project.

2 About CHIM

The main objective of CHIM is to develop a chatbot that can answer questions by museum visitors about objects in the museum. CHIM enables conversational interaction based on text and speech. Visitors can ask their questions and receive answers in multimodal formats (text, audio, image, video). In addition, the application will offer customized tours based on the interests and needs of the respective visitors to create a personalized experience.

In the process of developing CHIM, we explore different methods to extract answers from our corpus of existing audio guide texts. On the one hand, we explore how large language models, such as BERT [9], can be used in the museum chatbot context to find answers in unstructured or partially structured data. On the other hand, we explore how established methods for NLU can be efficiently integrated into the process of creating chatbot tours [7].

A crucial step in the creation of the CHIM chatbot is to build an NLU model for the classification of visitor questions. To collect relevant questions from potential museum visitors, we created a website designed specifically for this purpose. Our approach is to first identify the content types of the questions asked by museum visitors. To this end, we categorized the collected questions according to their content type. The question collection itself, the procedure for question categorization and the results of the categorization are described in the following section. In Sect. 4, we outline our planned and partially realized NLP strategies.

In a subsequent step, we will refine the content types by adding annotations for entities and relations. Further, to extract matching answers from the existing corpus of audio guide texts, the texts will also be labelled with content types.

3 Question Collection

3.1 Experimental Procedure

A website was built to gather relevant questions about 14 selected exhibits of the Städel Museum. To find as many contributors to the question collection as possible, a campaign was initiated in cooperation with the Städel Museum via the Städel Blog. In this way, we collected a total of 2182 questions from 203 unique user sessions during the period from December 22, 2020, to March 23, 2021. Each user session corresponds to one participant.

On the home page of the question collection website, we briefly described the procedure and purpose of the collection. The participants were presented a sub-selection of the 14 artworks, one at a time, and their task was to ask one or two questions per artwork. For each interaction, the date on which the interaction occurred, the input form (text input or voice input), as well as the browser used were anonymously stored. As input of the participants, the questions about the objects and optional comments about the application, as well as optional information about age, gender and education level were stored. About 50% of the participants provided demographic information. The average age of these participants was approximately 43 years (min. 17/max. 71). The participants had the following gender distribution: 63% female, 33% male, 2% other, 2% (explicitly) no indication. The educational background was distributed as follows: 80% university, 13% university of applied sciences, 5% high school, 2% other.

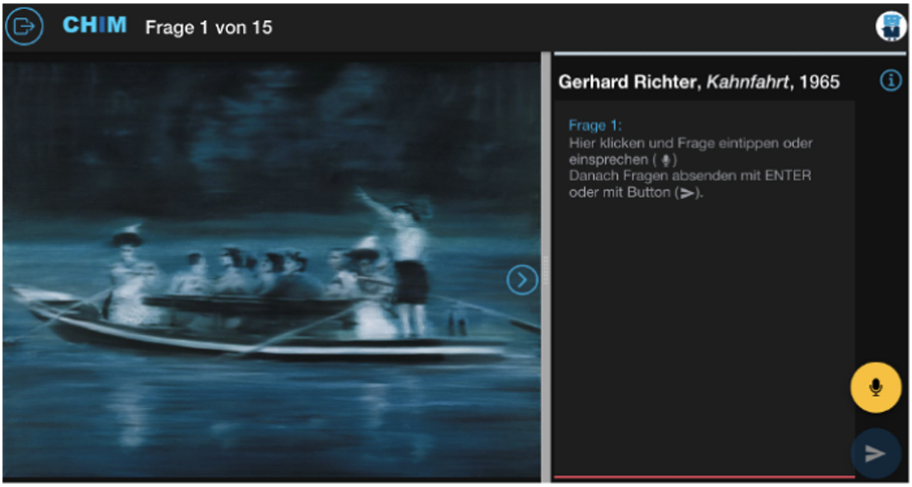

Fig. 1. Question Collection Website. Artwork is displayed in the Städel Museum. Photo: © Gerhard Richter 2020 (0217).

Figure 1 shows the user interface of the question collection website. Each participant was asked to enter a total of 15 questions. On the left side, below the question number, an image of the artwork was displayed for which questions were to be entered. At the top right, basic information like the artist's name, the title of the artwork and the year of creation were shown. Below this was a text field for entering the questions. Questions could be entered either via keyboard or by using the microphone symbol on the bottom right. Speech input was transcribed into text using automatic speech recognition. The recognized text was displayed in the text field. After entering 15 questions, a short questionnaire was displayed for demographic data and for comments about the application.

3.2 Question Categorization

The categorization of the questions is based on an approach of [6]. They categorized visitor questions into content types to explore which types of content voice-based AI conversational systems should attend to in order to meet visitors' expectations in a museum. We intend to utilize this approach to generate an annotated set of visitor questions that can be used to build an NLP model that is able to categorize questions about artworks into such content types. This model will be one of the building blocks in our NLP pipeline. The following 8 categories were used in [6]:

- *fact*: questions related to who is the artist, when the artwork was made, its size, or where it has been exhibited;
- *author*: visitor utterances about the artist's life, which art movement they were part of, or stylistic influences;
- *visual*: questions about colors and materials used, brushing techniques, etc.;
- *style*: questions about the style of the artwork, which school it belonged to and its characteristics, or artworks with style;
- *context*: inquiries about the historical, political, or social context where the artwork was produced;
- *meaning*: questions related to intentions, meanings, or whys, and the stories possibly behind the people and elements depicted in the artwork;
- *play*: utterances of playful engagement with the artwork, questions beyond the scope of the work, such as which soccer team a character roots for;
- *outside*: groups questions related to the conversational guide itself, its technology, or unrecognized utterances.

In their analysis, [6] revealed that far more than the half of the questions were about the *meaning* of the artworks (about 60%), followed by factual questions (17%), and questions about the artist's biography (7%). About 10% of the questions were not understood or were outside the scope of the artwork. The other 4 content types, together, corresponded to under 7% of the questions. Further, it was shown that the distribution of question types did not significantly differ per artwork.

As the content type *meaning* is overused in [6], we refinded this category by adding the following four content types:

- *content*: questions related to what or who is depicted in the artwork, both overall and in detail. Examples: "Is that the baby Jesus on her lap?", "Who are these people?", "Is the dog really sleeping or just pretending?"
- *model*: questions about the original models that were used. This is about the portrayed real person or object that have a real background. Examples: "Did the painter really work with a nursing mother as a model for this picture?", "Is the dog real?", "Surely this is painted from a photograph?"
- *response*: questions related to the response that the artwork triggers in/the effect the artwork has on the viewers, both historically and contemporary. Examples: "How did people back then react to the image of a bare breast?", "I think the picture is stupid, you can hardly see anything?", "What makes the painting so peaceful?"
- *provenance*: questions/information regarding the chronology of the ownership, custody, or location of the artwork. Examples: "How did this painting end up at the Städel Museum?", "Who commissioned the painting?", "Was this painting stolen by the Nazis?"

With the questions collected via our website, we ran a blind manual classification with five annotators. One main annotator created annotations for all questions, while the remaining four annotators annotated about 25% of the data each. Disagreements between the main and the other annotators were resolved jointly. When no consensus could be reached, the annotation of the main annotator was used. In this way, each question received exactly one annotation. In the next subsection, we will give an overview of the preliminary analyses of the annotated data.

3.3 Results and Discussion

Questions on All Artworks. Figure 2 illustrates the distribution of the questions we collected across the twelve content types in percentages. Approximately one third of the questions were about the *content* of the artworks (about 32%), followed by questions about the *meaning* (26.5%). The content types *visual*, *artist*, *model* and *context* all range between 5 and 9%. The content types *play*, *fact*, *style*, and *response* all range between 2 and 4%. Less than 2% of the questions were categorized as *outside* the scope of the artwork or as *provenance*.

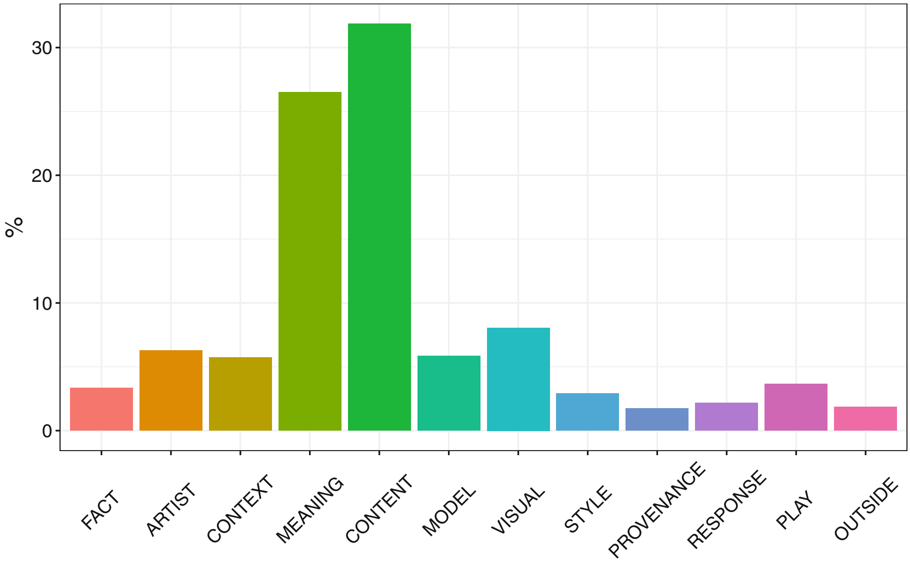

Fig. 2. Questions all artworks: distribution of questions according to content type in percentages. N = 2357.

Compared to [6], the content type *meaning* was considerably reduced from 60 to 26.5%. The largest contribution to this shift was made by the new content type *content*. The new content type *model* was the third most frequency, albeit contributing far less than the new content type *content* to the reduction of *meaning*. Out of the new content types, *response* and *provenance* were used the least. We conclude that adding more content types to split up the category *meaning* was successful in our case, since overall, a more balanced distribution of questions across the different content types was achieved.

The content type *fact* was used considerably less in our dataset than in [6]. This may be in part attributable to the inclusion of the new content type *provenance*, which can be seen as a subtype of *fact*. However, as *provenance* accounts for only 1.75% of the total questions, we also consider another explanation for this difference: The user interface of our question collection site already provides essential information for the category *fact* using text labels, displaying the artist's name, the title of the artwork and its year of creation. We deliberately chose this design, since in the Städel Museum, too, basic information is available on text labels displayed next to the artworks. However, it must be mentioned that as far as we know, the Pinacoteca museum in Brazil (the museum where the [6] application was tested) also displays basic information about the objects. We assume that sometimes this information is not easily visible for those visiting the exhibition. When collecting data in the future, we will consider not displaying such information on the website, so as to more closely mimic the actual situation in the exhibition.

Another clear difference is that the content type *outside* was used much less in our study. This can be explained by the fact that in [6], the category *outside* was used for annotation if the question was not understood by the system or was outside the domain of the artwork. In our study, so far no technical module is used to classify the questions, therefore, corresponding false detection in language understanding cannot occur.

Overall, the questions in our study are distributed more evenly across the content types than in [6]. In particular, our extension of the set by three additional content types may have contributed considerably to shift the distribution of the questions. A more balanced distribution is desirable for the creation of an NLP model: on the one hand, more training data is available for the classes of the model, avoiding biases due to uneven training data distribution. On the other hand, we hope that more clearly separated content types will lead to better precision determining the answers in further processing.

Looking into Data of Single Exhibits. Figure 3 shows a subset of our data. As is clearly visible, the distribution of the content types shows large differences for the individual images. For the object 'Lucca Madonna', the frequency of *context* and *meaning* is almost opposite of that of the overall distribution presented in Fig. 2. This painting from the field of Christian art is full of symbolic objects and imagery. We assume that this is one of the main reasons, why the questions are strongly concentrated on the meaning rather than the content.

Another notable difference can be seen with the object 'Boat Trip'. The content type *visual* is considerably more frequent than in the other distributions. Looking at the actual questions, we found that an above-average number of the questions relate to the technique the artist used to create the artwork.

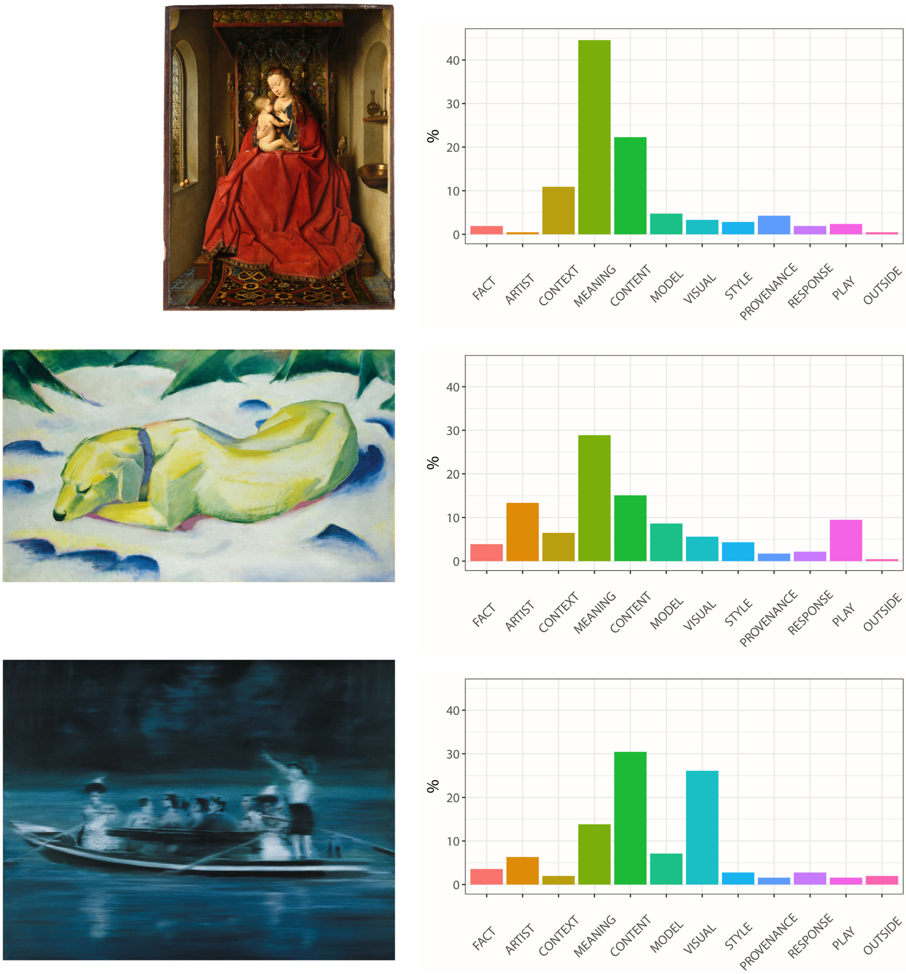

Fig. 3. Objects from top to bottom: Lucca Madonna[(a)], Boat Trip[(b)], Dog Lying in the Snow[(a)]; each with the distribution of questions across content types to the right. Photos: (a) CC BY-SA 4.0 Städel Museum, Frankfurt am Main.; (b) © Gerhard Richter 2020 (0217)

For the object 'Dog Lying in the Snow', increased usage of the content type *artist* can be observed. Again, looking at the actual questions, we found that many participants asked whether this was the artist's own pet dog, or if the artist liked to paint animals in general. Also noteworthy here is the use of the content type *play*. With this object, playful questions such as "does it bite?" were asked more frequently.

These preliminary results suggest a noticeable effect for the individual artworks on the frequency of specific content types in questions. However, this contrasts with the results of [6]. They found no significant correlation between artwork and content type. One possible explanation is that for the participants of our survey, each artwork represents a domain of its own. Across different domain, the content types may differ. However, this is only a preliminary finding. So far, we have not been able to extensively investigate the data of all the works. Going forward in our project we plan to investigate this difference and possible explanations for it.

4 Answering Strategies

The main task of the chatbot is to give a satisfactory answer to users' questions. This can be framed within the classical NLP problem of Question Answering (QA), i.e., based on a question, finding the correct document or excerpt within a document that contains the answer.

For the documents containing the answers, we considered two options: the first is to create dedicated answers specifically designed (= written) for the chatbot, whereas the second is to utilize existing text documents and descriptions for the exhibits in question - in the case of our project, the corpus of existing audio guide texts. When using existing texts, different degrees of enriching the text with metadata are possible (see Table 1), that allow better "machine understanding".

For example, the sentence "His way of painting was radically different from the International Gothic style, which at that time had been prevalent across Europe." could be annotated with metadata "artist: Jan van Eyck" making the artist in this sentence explicit; or annotated with some metadata like "style" as a content type to indicate that this sentence deals with the style of a painting.

While creating a specific answer for each question would be ideal, it is also the costliest option regarding time and effort. In addition, these dedicated answers can only cover those questions, or answers to those questions, that occur in the corpus of collected questions. Topics that are not brought up by these questions will in principle not be answerable by dedicated answers. In this event, the existing audio guide texts can be used as a fallback, since these texts usually are written with the goal to cover a wide variety of informational needs.

4.1 Degree of Enrichment

The effort for utilizing existing text depends on the degree of "enrichment". In our project, we follow a multi-tiered approach where we apply different degrees of enrichment and effort for the answers: for a few selected exhibits, we will create new, dedicated answers as well as highly metadata-enriched texts. For the rest of the exhibits, only question-clusters that crystallized as "frequently asked" by different users during our annotation phase will get dedicated written answers, and only, if these answers do not already exist in the available texts. Furthermore, these remaining exhibits' text descriptions will receive a middle to low degree of effort regarding metadata enrichment. One goal in our project is to find out exactly which degree of effort is minimally necessary or is enough to be able to create a satisfiable user experience.

When designing dedicated answers, it is useful to classify questions as factoid-type and open-ended-type questions: generally, factoid-type questions aim at short, specific answers (e.g., "when was this painted?") while open-ended-type questions require more elaboration (e.g., "what is the meaning of the painting"). Preliminary evaluations have shown that depending on the type of question, different NLP approaches are more - or less - successful at delivering satisfying answers [14, 15].

Accordingly, our system uses multiple NLP techniques in stages, depending on how likely they are to deliver a satisfying answer. The mechanisms are designed so as to always return some kind of relevant information; if fallback mechanisms (see Sect. 4.2) are used, the system explicitly states that it may not have found *the* answer, but some information potentially related to the requested answer. In addition, the chatbot will employ further interaction strategies for dealing with errors and failures, which are not discussed in this paper.

Table 1. Used NLP mechanisms and their required vs. optional metadata-enrichments. Abbreviations: Entity (E), Relationship (R), Event (Ev), Content Type (CT).

	Required enrichment	Optional/additional enrichment
Factoid-ER module	E/R/Ev annotation in (set of) questions	
Factoid-BERT module		CT ("intent") annotations (in answer sentences)
Open-ended intent module	CT ("intent") annotations (in answer sentences)	E/R/Ev annotations (in questions & answers)
Open-ended similarity module	"section annotations": annotated answer sentences w.r.t. "same topic"/coherence Text follows	CT ("intent") annotation of answer sentences; E/R/Ev annotations (in questions & answers)

4.2 NLP Techniques

During the first stage, Intent Recognition using the tool Rasa [7] trained on the content type annotations is applied as well as a classification for factoid- or open-ended-type of question. For factoid-type questions, an Entity Relation Extraction mechanism will try to identify the question-target and -topic (e.g., "when was the image painted?": target is image, topic is time-of-creation). If successful, the corresponding factoid-datum is retrieved from a database and a natural language answer is generated. If unsuccessful, a BERT [9] model, pre-trained for QA is utilized for finding a matching answer in the text documents available for that particular exhibit.

For open-ended-type questions, answer candidates are retrieved from the dedicated answers and the annotated audio guide texts, if their annotated content type matches the recognized content type of the user's question with sufficiently high confidence. If the answer candidates comprise a continuous section of text, this longer explanation will be selected as answer. If the answer candidates correspond to multiple, separate sections, we plan to use Entity Extraction to reduce answer-candidates further down.

If the confidence for recognizing the question's intent is not sufficiently high to extract answer candidates based on this feature, a fallback mechanism is used that calculates a cosine-similarity [12] between the question and all answer-sentences, and then selects the encompassing text-section of the sentence with the highest similarity. From this, a chatbot answer is created, stating that no matching document could be found, but maybe the returned text contains some related information.

5 Conclusion and Future Work

In this paper, we report essential steps that we undertook in the development of CHIM, the prototype of a chatbot AI that answers questions of visitors in the museum. We adapted an existing approach for categorizing questions of potential museum visitors according to content type. In our adapted approach, we increased the number of content type categories from eight to eleven and used them to categorize a set of questions about exhibits in the Städel Museum. In the original approach, the distribution heavily skewed in favor of the content type *meaning*. Preliminary results from our extended approach show a more balanced distribution of the questions across the different content types. The annotated questions are used to set up a multi-tiered NLP approach in which we apply different degrees of effort to generate answers. Compared to building a complex NLP model using a multitude of fine-grained intents and entities, categorizing questions into rather rough content types comprises a comparably low degree of metadata enrichment and thus less annotation work for human annotators. To ensure scalability during the production of museum chatbots that make use of existing audio guide content, we suggest making use of highly metadata-enriched texts only for a few selected exhibits. For the rest of the exhibits, we reduce the overall effort for enrichment by applying NLP mechanisms trained with content type annotated questions and a follow-up NLP mechanism like BERT that can be applied to unstructured data. By using this approach, we consider that there are finite resources for creating "chatbot content", balancing the effort of creating new content and making existing content usable by enriching it to varying degrees. Our future work in the project CHIM will include the implementation of dialog management, a further refinement of the NLP strategies, and the integration of additionally created answers for selected exhibits. Further results will be reported after the final implementation and evaluation of the system during a field test at the Städel Museum in late 2021.

References

1. The field museum. https://www.fieldmuseum.org/exhibitions/maximo-titanosaur?chat=open. Accessed 30 July 2021
2. Jüdisches museum berlin. https://www.jmberlin.de/whatsapp-guide-hey-und-herzlich-willkommen. Accessed 30 July 2021
3. Kunsthalle karlsruhe: Art of chit-chatting. https://www.moodfor.art/chit-chatting. Accessed 02 Nov 2021
4. Ping! die museumsapp. https://www.museum4punkt0.de/ergebnis/ping-die-museumsapp-spielerisch-durchs-museum. Accessed 02 Nov 2021
5. Zentrum für kunst und medien. https://zkm.de/de/talk-to-me-chatbots-in-museen. Accessed 30 July 2021
6. Barth, F., Candello, H., Cavalin, P., Pinhanez, C.: Intentions, meanings, and whys: designing content for voice-based conversational museum guides. In: Proceedings of the 2nd Conference on Conversational User Interfaces, pp. 1–8 (2020)
7. Bocklisch, T., Faulkner, J., Pawlowski, N., Nichol, A.: Rasa: open source language understanding and dialogue management. arXiv preprint arXiv:1712.05181 (2017)
8. Ciecko, B.: Examining the impact of artificial intelligence in museums, February 2017
9. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. CoRR abs/1810.04805 (2018). http://arxiv.org/abs/1810.04805
10. Falk, J., Dierking, L.: Learning from museums: visitor experiences and the making of meaning, January 2000
11. Gaia, G., Boiano, S., Borda, A.: Engaging museum visitors with AI: the case of chatbots. In: Giannini, T., Bowen, J.P. (eds.) Museums and Digital Culture. SSCC, pp. 309–329. Springer, Cham (2019). https://doi.org/10.1007/978-3-319-97457-6_15
12. Huang, A.: Similarity measures for text document clustering. In: Proceedings of the Sixth New Zealand Computer Science Research Student Conference (NZCSRSC 2008), Christchurch, New Zealand, vol. 4, pp. 9–56 (2008)
13. Kohle, H.: Digitale Bildwissenschaft. Hülsbusch, Glückstadt (2013). http://nbn-resolving.de/urn/resolver.pl?urn=nbn:de:bvb:19-epub-25747-3
14. Zaman, M.M.U., Schaffer, S., Scheffler, T.: Comparing BERT with an intent based question answering setup for open-ended questions in the museum domain. In: 32. Konferenz Elektronische Sprachsignalverarbeitung. Elektronische Sprachsignalverarbeitung. Elektronische Sprachsignalverarbeitung (ESSV-2021). TUDpress, Dresden (2021)
15. Zaman, M.M.U., Schaffer, S., Scheffler, T.: Factoid and open-ended question answering with BERT in the museum domain. In: Proceedings of the Conference on Digital Curation Technologies. Conference on Digital Curation Technologies (QURATOR-2021). CEUR Workshop Proceedings (2021)