



# Is GPT-3 All You Need for Visual Question Answering in Cultural Heritage?

Pietro Bongini<sup>✉</sup>, Federico Becattini<sup>✉</sup>, and Alberto Del Bimbo<sup>✉</sup>

University of Florence, MICC, Florence, Italy  
{Pietro.Bongini,Federico.Becattini,AlbertoDel.Bimbo}@unifi.it

**Abstract.** The use of Deep Learning and Computer Vision in the Cultural Heritage domain is becoming highly relevant in the last few years with lots of applications about audio smart guides, interactive museums and augmented reality. All these technologies require lots of data to work effectively and be useful for the user. In the context of artworks, such data is annotated by experts in an expensive and time consuming process. In particular, for each artwork, an image of the artwork and a description sheet have to be collected in order to perform common tasks like Visual Question Answering. In this paper we propose a method for Visual Question Answering that allows to generate at runtime a description sheet that can be used for answering both visual and contextual questions about the artwork, avoiding completely the image and the annotation process. For this purpose, we investigate on the use of GPT-3 for generating descriptions for artworks analyzing the quality of generated descriptions through captioning metrics. Finally we evaluate the performance for Visual Question Answering and captioning tasks.

**Keywords:** Visual Question Answering · GPT-3 · Image captioning · Natural language processing · Computer vision

## 1 Introduction

Cultural Heritage often relies on digital resources to engage and attract visitors. From audio-guides to smartphone applications, museum visits are becoming increasingly more interactive, allowing users to deepen concepts without the need of a human assistant or after the visit is concluded. Forms of gamification are also important, favoring engagement especially for young visitors and instructional purposes. Artificial Intelligence and Computer Vision are playing a large part in the development of such smart visits and applications [5, 6, 8, 13]. A notable machine learning application that has recently found usage in cultural heritage is Visual Question Answering (VQA), which exploits both Computer Vision and Natural Language Processing to allow users to ask questions on the content of an image [6]. The advantage of VQA is that it allows museums to develop smart guides and interactive gamification approaches. However, for pictorial art, most questions posed by users concern contextual information rather than what is actually depicted in a painting.

To address this limitation, an evolution of VQA known as Contextual Question Answering (CQA) was proposed [6]. The authors explicitly focused on cultural heritage applications, combining visual and contextual cues to answer questions. The contextual information is derived from a textual meta-data, which is fed to the model along with the question and the image. In this way the VQA/CQA model has to learn to attend either relevant parts of an image or relevant sections of the text to provide an adequate answer. The need of a textual data nonetheless opens a new issue, namely where to obtain such description. Information sheets for artworks may already be available to museum curators yet extending this kind of application to new data becomes time-consuming and requires a domain expert.

In this paper we explore the usage of a generative natural language processing model to automatically create contextual information to be fed to a CQA model. In fact, recently, generative text models have been finding large diffusion with groundbreaking results. Among these we find GPT-3, a generative model trained on a massive corpus of textual data regarding several domains, including art [7]. GPT-3 is capable of generating a description starting from a textual query and it has been demonstrated that the model includes knowledge of the entities described in the training data, for example paintings and artworks. We therefore investigate the possibilities and the limitations of GPT-3 in applications for cultural heritage, with a specific focus on question answering. In particular, we explore the quality of the textual description of artworks that the model is able to generate and we evaluate their applicability for visual and contextual question answering.

The main contributions of our work are the following:

- We propose an automatic approach to generate textual information sheets of artworks exploiting GPT-3. We find that the model has excellent knowledge of art concepts and event details of specific paintings.
- We propose a method to answer both visual and contextual questions which is artwork agnostic, i.e. it does not require any additional data or training to be adapted to a new set of images.
- We explore the applicability of GPT-3 in cultural heritage applications. To the best of our knowledge we are the first to apply GPT-3 to the art domain.

## 2 Related Work

Natural Language Processing (NLP) in recent years has evolved at an extremely fast pace, converging to a set of well defined application paradigms [33]. Such paradigms include text classification, matching, machine reading comprehension, sequence to sequence translation, sequence tagging and language modeling. Despite the wide variety of tasks [1, 9, 28], some recent noticeable approaches have been shown to perform well as generic pre-training for NLP models [7, 11]. In particular, this can be attributed to the introduction of attention models, based on the transformer architecture [36]. The effectiveness of models such as BERT [11] stems from the capability of processing text bidirectionally exploiting the

self-attention mechanism of transformers to obtain word level representations that are informed of their surrounding context within the sentence. Whereas BERT is built exploiting the encoder part of the transformers, another state of the art approach for NLP, Generative Pre-trained Transformer (GPT) [22], is built stacking transformer decoder blocks and is trained to predict the next word in a sentence. The model has then been improved in subsequent versions, GPT-2 [23] and GPT-3 [7], yielding larger and more effective models.

Interestingly, GPT-3 has been trained using a large quantity of internet data, meaning that the training process has distilled into the model common sense knowledge making it able to generate essays and even poetry [10]. In this paper we exploit GPT-3 as a generator of textual content describing artworks, showing that it can be used for interactive applications for cultural heritage such as captioning [19] and Visual Question Answering (VQA) [2]. VQA is a recent trend in machine learning that bridges the Natural Language Processing and Computer Vision domains [4]. The goal is to answer questions regarding the content of an image through artificial intelligence. This involves several sub-tasks such as object detection [15] and recognition [16], question reasoning [20]. Typical VQA approaches use Convolutional Neural Networks (CNNs) to interpret images and Recurrent Neural Networks (RNNs) to process questions. The authors of [1] proposed a bottom-up attention mechanism looking at salient objects in images. Differently from previous approaches that considered regularly spaced image portions [30], they use object Faster R-CNN [25] features as attention candidates. In the past few years multiple Transformer-based approaches reached impressive performances on this task [17, 32, 34, 38].

Recently, a few approaches [3, 6, 14, 35] have addressed VQA in the cultural heritage domain. A dataset of questions and answers for art related questions has been recently proposed [3], exploiting an ontology based framework to extract data with question templates. The authors of [6] and [14] found that to make the best out of VQA for museum applications, a model must be able to integrate some source of external knowledge in order to address contextual questions, i.e. questions concerning non-visual cues such as name of the author, year and artistic style. In particular, [6] used a question classifier to understand if visual or contextual knowledge is required. Depending on the output of the classifier a VQA model is used, otherwise a purely textual based question answering model is used discarding the image content. In this work we explore the effectiveness of using GPT-3 to generate artwork captions, suitable for such a visual and contextual question answering model.

Other approaches have been used to answer questions relying on captions, yet only regarding visual content [29]. The most similar approach to ours is instead [39], which used GPT-3 for VQA. However, differently from us, the authors feed GPT-3 with questions and descriptions generated by an image captioner directly to obtain an answer. We, instead, aim at extracting the domain specific knowledge from GPT-3 which is requested to correctly answer a question.

### 3 GPT-3

To provide to the reader a better understanding of our work, here we present a brief background context about GPT-3, the third version of Generative Pre-Trained Transformer [7]. This is an autoregressive language model with 175 billion parameters that can be used for different tasks without any finetuning, achieving strong performances.

The architecture of the GPT-3 Transformer model is made of 96 attention layers. While language models like BERT [11] use the Encoder to generate embeddings from the raw text which can be used in other machine learning applications, GPT-3 use the Decoder half, so it takes embeddings as inputs and produces text. In particular the GPT-3 language model has the ability to generate natural language text that can be hard to distinguish from human-written text, to the point that research has been carried out to asses whether GPT-3 could pass a written Turing test [12].

Concretely, during inference, the target of the new task  $y$  is directly predicted conditioned on the given context  $C$  and the new task’s input  $x$ , as a text sequence generation task. Note that all  $C$ ,  $x$  and  $y$  are text sequences. For example,  $y = (y^1, \dots, y^T)$ . Therefore, at each decoding step  $t$  we have

$$y^t = \arg \max_{y^t} p_W(y^t | C, x, y < t) \quad (1)$$

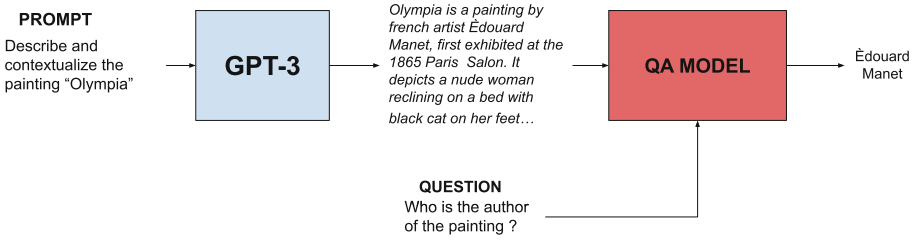
where  $W$  are the weights of the pretrained language model, which are frozen for all new tasks. The context  $C = h, x_1, y_1, \dots, x_n, y_n$  consists of an optional prompt head  $h$  and  $n$  in-context examples  $(\{x_i, y_i\}_{i=1}^n)$  from the new task.

### 4 Method

In a Cultural Heritage context, the information useful to answer questions about a specific artwork is contained in the artwork image and in its contextual description. Finding such a description might not be trivial, since it might require a domain expert to write it down. At the same time, it is quite costly to train a Visual Question Answering model that takes in input both the image and the description. This is also not straightforward, since the two modalities need to be blended and matched together. Consequently, the main idea of this work is to generate new descriptions for artworks based on a specific prompt or a specific question and directly use these descriptions to answer visual and contextual questions. The overall pipeline of our proposed work is as follows:

1. **GPT-3 caption generation.** We use GPT-3 to generate descriptions of artworks, leveraging its memorization capabilities that allowed the model retain relevant information about training instances. An important aspect in this phase in to feed the correct prompt in input to GPT-3 in order to obtain realistic and correct descriptions. We consider two different types of input prompt:

**General description generation for question answering**




**Fig. 1.** Scheme of our method for answering questions using a general generated description. A prompt with a specific structure is given in input to GPT-3. Subsequently the generated text is fed together with the question to a Question Answering model that outputs the answer.

- **General** - A general prompt where the expected output is a general description of the artwork. The input text follows the structure:  
 "Describe and Contextualize the painting < painting\_name >"
- **Question-based** - A specific question based prompt. The input text follows the structure:  
 "Painting < painting\_name > < question >".  
 The expected generated text by GPT-3 is a small text snippet that consists in a couple of sentences, focused on the topic of the question.

2. **Question answering.** Once the description has been generated in the previous step, we can exploit it to answer both visual and contextual questions through a Question Answering language model. For this purpose we use a pretrained version of DistilBert [27] fine-tuned on the SQUAD [24] dataset. We feed in input to the DistilBert model the generated text from the previous step together with the question. The answer given as output will be the final answer of our method.

Figures 1 and 2 show a scheme of the two variants of our method. More precisely, in Fig. 1 the general input prompt for GPT-3 yields the generation of a long description of the artwork (similar to a museum information sheet). On the other hand, the question-based prompt in Fig. 2 yields only the generation of a brief output text, which we find suitable for answering the question. In conclusion, these two schemes follow roughly the same structure. The difference is in the input prompt that in the case of Fig. 1 is more general and in Fig. 2 is more task oriented.

**Question-based description generation for question answering**

**Fig. 2.** Scheme of our method for answering questions using a question-based generated description. A prompt containing the name of the painting and the question is given in input to GPT-3. Subsequently the generated text is fed together with the question to a Question Answering model that outputs the answer.

## 5 Experiments

In this section we first outline the experimental setting for the experiments carried out in this paper, presenting dataset and experimental protocol and we then move on to a discussion of the results.

### 5.1 Dataset

For our experiments, we use the Artpedia dataset [31]. Artpedia contains a collection of 2,930 artworks, associated to a variable number of textual descriptions gathered from Wikipedia. Sentences are labelled as a visual descriptions or as a contextual descriptions. Contextual descriptions regard information about the artwork that does not directly describe its visual content. For instance, contextual descriptions can describe the historical context of the artwork, its author, the artistic influence or the museum where a painting is exhibited. The dataset contains 28,212 descriptions, 9,173 of which are labelled as visual and the remaining 19,039 as contextual. The Artpedia dataset has been extended with Question-Answer annotations in [6]. In fact, a subset of the images have been associated with visual and contextual questions, derived from the corresponding captions. In this work we follow the dataset split of [6].

### 5.2 Experimental Protocol

Following prior work such as [6], we evaluate visual questions and contextual questions with different metrics. In fact, visual question answering and traditional text-based question answering are often treated in two different ways. Visual Question Answering is considered as a classification problem, meaning that a model has to pick an answer from a predefined dictionary of possible candidates containing a few words each. This stems from the fact that questions in most datasets are a way of guiding attention towards specific objects

or attributes in the image, without requiring any complex form of language reasoning. Question Answering on the other hand is based on a set of sentences, which may contain rare or out-of-dictionary words. The task is in fact defined as identifying a subset of the textual description that contains the answer.

In light of this, to evaluate visual questions we rely on accuracy:

$$Accuracy = \frac{N_c}{N_a} \quad (2)$$

where  $N_c$  is the number of correct answers and  $N_a$  the number of total answers.

For text-based question answering, instead, we use both accuracy and F1-measure, a metric that takes into account the global correctness of the answer:

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (3)$$

Where *Precision* is defined as:

$$Precision = \frac{N_{Cw}}{|ans|} \quad (4)$$

with  $N_{Cw}$  is the number of common words between the output answer and the ground truth answer and *ans* the number of words in the generated answer.

*Recall* instead is defined as:

$$Recall = \frac{N_{Cw}}{|gt|} \quad (5)$$

where  $|gt|$  is the number of words in the ground truth.

We also evaluate the quality of the descriptions generated by GPT-3, considering it as a standalone image captioning model. We use the following standard metrics for captioning:

- *BLEU1* [21]: BiLingual Evaluation Understudy (BLEU) is the most commonly used metric for machine translation and image captioning. BLEU scores are based on how similar a generated caption is to a reference caption, computing the precision of the generated words. The downside of BLEU is that it is very sensitive to small changes, such as synonyms or different word order.
- *ROUGE* [18]: differently from BLEU, which measures the precision of the caption, Recall Oriented Understudy of Gisting Evaluation (ROUGE) focuses on quantifying the amount of correct words with respect to the reference. Thus, this metric is recall-based and tends to reward long sentences.
- *CIDEr* [37]: Consensus-based Image Description Evaluation (CIDEr) is an automatic consensus metric that measures the similarity of captions against a set of ground truth sentences written by humans. This metric has been shown to yield a higher agreement with humans generated text since it captures notions of grammar, importance and precision and recall.

**Table 1.** Image captioning results. We compare our method which generates captions with GPT-3 with the *General* and the *Question-based* approaches. In the *Question-based* approach we concatenate all the outputs of GPT-3 after conditioning it with different questions related to the image. We compare the results against visual captions, contextual captions or both.

Description type	Metric	OFA [38]	Ours General	Ours Question-based
Visual	BLEU1	0.048	<b>0.181</b>	0.137
	ROUGE	0.138	<b>0.188</b>	0.16
	CIDEr	0.091	0.079	<b>0.172</b>
	COSINE	0.113	<b>0.157</b>	0.110
Contextual	BLEU1	0.002	<b>0.168</b>	0.160
	ROUGE	0.062	0.178	<b>0.179</b>
	CIDEr	0.000	<b>0.248</b>	0.129
	COSINE	0.082	0.218	<b>0.324</b>
All	BLEU1	0.000	0.113	<b>0.185</b>
	ROUGE	0.053	0.158	<b>0.184</b>
	CIDEr	0.000	0.016	<b>0.098</b>
	COSINE	0.122	0.253	<b>0.341</b>

- *Cosine Similarity*: we compute the cosine similarity between feature vectors for the generated caption and the reference caption. Features are extracted with the algorithm TF-IDF [26].

### 5.3 Experimental Results

**Captioning Results.** We start by assessing the quality of the captions generated by GPT-3. First of all, we ask GPT-3 to generate captions with our *General* approach. In Table 1 we compare the captions using as reference visual captions, contextual captions and both. All reference captions are ground truth captions taken from the Artpedia dataset [31].

Interestingly, the model appears to better results for visual captions using BLEU1 and ROUGE metrics, while using CIDEr and cosine similarity, the model obtains higher results for contextual captions. This may seem counter-intuitive but can be explained looking at the nature of the metrics. BLEU1 and ROUGE in fact respectively check for word-wise precision and recall, while CIDEr and cosine distance perform a sentence level scoring, which is closer to human consensus. We observe that the model is able to obtain good results, especially with the cosine metric, even when using all the captions as reference.

We then evaluate the method by taking a concatenation of the outputs generated by GPT-3 after being conditioned by different questions related to the image. This obviously introduces a strong bias, given also the fact that questions have been generated from information contained in the captions, but at the same time proves the usefulness of such captions for more advanced applications such as visual question answering. As can be seen in Table 1, conditioning GPT-3 with the captions leads to better captions according to most metrics.



**Table 2.** Experimental results for Visual Question Answering. We compare our approach against VQA-CH [6] to understand whether GPT-3 can replace information sheets for artworks either for visual or contextual questions. We compare two versions of our model, the *General* version, which produces generic descriptions of artworks and the *Question-based* version, where prompts are conditioned with the input question to generate more specific descriptions.

	Visual	Contextual	Accuracy	F1 score
VQA-CH [6]	✗	✓	0.684	0.832
VQA-CH [6]	✓	✗	0.176	0.150
VQA-CH [6]	✓	✓	0.504	0.417
Ours - General	✗	✓	0.557	0.719
Ours - General	✓	✗	0.070	0.055
Ours - General	✓	✓	0.239	0.360
Ours - Question-based	✗	✓	0.473	0.602
Ours - Question-based	✓	✗	0.134	0.202
Ours - Question-based	✓	✓	0.256	0.330

In Table 1 we also provide a baseline as reference, i.e. the output of the state of the art OFA captioning model [38]. We observe that captions generated by OFA do not align well with the ground truth sentences. We attribute this to a domain shift between the datasets commonly used to train captioning models and descriptions of artworks. In fact, the former are sentences written by non-experts while for applications in cultural heritage a domain knowledge is required. This further motivates the usage of GPT-3, which seems to have integrated sufficient knowledge to articulate complex sentences with a domain specific jargon.

**VQA Results.** To evaluate the Visual Question Answering capabilities of our proposed method, we follow the setting of [6]. However, we do not rely on any vision-based model but rather on a fully textual question answering model based on DistilBert [27], as explained in Sect. 4. In Table 2, we compare our approach to the one of VQA-CH [6]. It has to be noted that, contrary to [6], we do not rely on real textual descriptions, which are known to contain the answer, but we only extract information from GPT-3. This is a strong disadvantage for our method. However, we are not interested in obtaining better results than VQA-CH, but rather our goal is to demonstrate if GPT-3 can act as a substitute of textual descriptions handcrafted by domain experts.

We test our method evaluating the accuracy for visual questions, contextual questions and both together. Quantitative results indicate that captions generated by GPT-3 can yield to high results for contextual questions, yet very low accuracy for visual questions. As for the captioning setting, we impute this behavior to the fact that GPT-3 generates generic descriptions, without including a fine-grained description of the visual content. Thus, on the one hand the question answering model is capable of extracting meaningful information from the generated captions. This means that GPT-3 is indeed capable of integrating domain knowledge during training and is capable of generating a complete

ARTEPIDIA DESCRIPTION	QUESTION	GENERAL GENERATED DESCRIPTION	GENERAL ANSWER	QUESTION-BASED GENERATED DESCRIPTION	QUESTION-BASED ANSWER
<p><i>Bathers at Asnières (French: Une Baignade, Asnières)</i> is an oil-on-canvas painting by the French artist Georges Pierre Seurat, the first of his two masterpieces on the monumental scale. "The canvas is of a suburban, placid Parisian riverside scene." A combination of complex brushstroke techniques, and a meticulous application of contemporary colour theory bring to the composition a sense of gentle...</p>	Who is the author of this painting?	<p><i>Bathers at Asnières</i> is a painting by French artist Georges Seurat. The painting depicts a group of men relaxing by the river Seine on a sunny day. The background includes industrial buildings and smokestacks, which contrast with the leisurely scene in the foreground. Seurat used small dots of color to create a sense of light and movement in the painting, which has been called an early example of Pointillism. The painting was completed in 1884, just before Seurat's death at the age of 31. It was exhibited posthumously at the Salon des Indépendants in 1886, where it received mixed reviews from critics. Some praised its technical mastery, while others found it cold</p>	Georges Seurat ✓	The painting <i>Bathers at Asnières</i> was created by french artist George Seurat in 1884	Georges Seurat ✓
	When was the painting depicted?		1884 ✓	<i>Bathers at Asnières</i> was painted by George Seurat in 1884	1884 ✓
	What is the animal in the ground?		✗	The animal on the ground is a dog	dog ✓
	Is there a man lying down?		✗	There is a man lying down in the painting <i>Bathers at Asnières</i> , but it is difficult to see. He is on the left side of the painting, near the bottom.	There is a man lying down ✓
<p><i>Dogs Playing Poker</i>, by Cassius Marcellus Coolidge, refers collectively to an 1894 painting, a 1903 series of sixteen oil paintings commissioned by Brown &amp; Bigelow to advertise cigars, and a 1910 painting. "Critic Annette Ferrara has described <i>Dogs Playing Poker</i> as "indelibly burned into... All eighteen paintings in the overall series feature anthropomorphized dogs, but the eleven in which dogs are..."</p>	Who is the author of this painting?	<p><i>Dogs Playing Poker</i> is a painting by American artist Cassius Marcellus Coolidge. The painting depicts a group of dogs sitting around a table playing poker. The dogs are all wearing human clothes and have cigars in their mouths. The painting is set in a dimly lit room, with a fireplace in the background. The painting is often seen as an example of kitsch art, due to its subject matter and style. It is also seen as a classic example of Americana, due to its depiction of a group of dogs engaged in a typically American activity.</p>	Cassius Marcellus Coolidge ✓	The painting is anonymous.	✗
	Who is sitting on the chair?		dogs ✓	The artist who painted <i>Dogs Playing Poker</i> is named C. M. Coolidge.	✗
	What are the dogs doing?		playing poker ✓	The dogs in the painting are playing poker.	dog ✓
	What is the color of the table?		✗	table in <i>Dogs Playing Poker</i> is green	green ✓
<p><i>The Singing Butler</i> is an oil-on-canvas painting made by Scottish artist Jack Vettriano in 1992. As a contemporary cultural icon, <i>The Singing Butler</i> has been compared to Grant Wood's <i>American Gothic</i>. "It depicts a couple dancing on the damp sand of a beach on the coast of Fife, with grey skies above a low horizon. To the left and right, a maid and a man hold up umbrellas against the weather..."</p>	When was the painting depicted?	<p>The painting "<i>The Singing Butler</i>" was painted by Scottish artist Jack Vettriano in 1992. The painting depicts two people, a man and a woman, standing on a beach with a butler who is singing and playing the guitar. The background of the painting is a blue sky with white clouds. The painting is set in the early 20th century.</p>	1992 ✓	This painting was depicted in 1992.	1992 ✓
	What are the two people in the middle doing?		dancing ✓	The two people in the middle are dancing	dancing ✓
	How many umbrellas are there?		✗	There are two umbrellas in the painting.	two ✓

**Fig. 3.** Qualitative results of our method. *Green*: ground truth description from the Artpedia dataset [31] and input question. *Yellow*: general descriptions provided by GPT-3 and answer obtained based on such text. *Blue*: Question-based description and correspondent answer. General descriptions are longer and more detailed than question-based generated descriptions. However, question-based generated descriptions are customized for the specific question. (Color figure online)

information sheet of the artwork. On the other hand, captions appear to be too generic to obtain information about specific details in the image.

To overcome this limitation, we test the model using captions generated with out *Question-based* approach. By feeding the answer to GPT-3 along with the title of the artwork, the model is able to generate more specific captions. Such captions, as explained in Sect. 6 are usually shorter but are focused on the prompt. This is particularly interesting since it means that a purely text-based model is capable of addressing a vision-based task. In Table 2 it can be seen that for visual questions alone, our method with question-based captions performs on par or better than the vision-based VQA-CH model.

## 6 Qualitative Analysis

In this section we provide a qualitative analysis of the captions generated by GPT-3 in order to characterize which kind of information they contain in both the *General* and *Question-based* formulation.

Since the prompts that we feed to GPT-3 are different, with one being more general and the other being question-based, we expect that the corresponding generated text by GPT-3 will be different. In Fig. 3 we can observe these differences. Generated general descriptions are very long and have the aspect of artwork information sheets in which we can find some visual and contextual

information. Question-based generated descriptions are instead shorter and contain the knowledge needed to answer to the specific questions. From Fig. 3 we can observe that the general description is very useful to answer to contextual questions but fails on some visual questions. This is likely due to different reasons:

- The generated text does not take into account any specific question and this can lead to the generation of a description without specific information useful to answer to the question.
- Visual questions are very specific since they refer to object relationships, colors, counting, etc. and the GPT-3 model tends to be more shallow in generating its descriptions.

On the other hand, question-based generated descriptions are helpful to answer visual questions but the small generated description useful to answer those specific questions could contain incorrect information leading to wrong answer predictions. In conclusion these two ways of generating text to answer visual and contextual questions have some pros and cons:

- General descriptions are longer and contain several pieces of information about the artwork. However this is fixed and could not contain the information needed to answer some questions.
- Question-based descriptions are generated for specific questions and contain only the information needed to answer the question on which GPT-3 has been conditioned. If the model has not memorized any specific information regarding such questions it may contain mistakes and descriptions will have to be re-computed for each question.

## 7 Considerations on Complexity and Accessibility of GPT-3

In the previous sections we have demonstrated that GPT-3 could indeed replace the usage of an information sheet handcrafted by a domain expert. However, we need to understand the actual applicability of GPT-3 in a real case application. GPT-3 has 175B parameters, which approximately amounts to 700GB. This means that inference on a single GPU is unfeasible due to current technological limits. The model however has been made available from OpenAI and is accessible through API that have a pricing fee per generated token. These considerations somewhat limit a large-scale usage of the model, especially if a description has to be generated for each question to be answered. On the other hand, generating fixed descriptions offline, one for each artwork, appears a viable solution at least for addressing contextual questions.

## 8 Conclusions

In this paper we presented a method for Visual Question Answering in the Cultural Heritage domain. In particular we have addressed the problem of data

annotation for artworks, generating descriptions with GPT-3. The performances for the VQA task show that the generated descriptions are useful to answer the questions correctly. This technique allows to answer visual and contextual questions focusing only on the generated description and can be used for any artwork. In fact, there is no need to retrain the model to incorporate new knowledge. This is possible thanks to the memorization capabilities of GPT-3, which at training time has observed millions of tokens regarding domain-specific knowledge. Finally the generated description can be integrated as textual input (textual description) in a more complex architecture as [6] in order to address tasks like Visual Question Answering. This is of particular interest for Cultural Heritage due to the domain shift between common VQA and captioning datasets compared to the technical jargon that is needed to properly address questions about art.

**Acknowledgement.** This work was partially supported by the project ARS01 00421: “PON IDEHA - Innovazioni per l’elaborazione dei dati nel settore del Patrimonio Culturale.” This work was partially supported by the European Commission under European Horizon 2020 Programme, grant number 101004545 - ReInHerit.

## References

1. Anderson, P., et al.: Bottom-up and top-down attention for image captioning and visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6077–6086 (2018)
2. Antol, S., et al.: VQA: visual question answering. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2425–2433 (2015)
3. Asprino, L., Bulla, L., Marinucci, L., Mongiovì, M., Presutti, V.: A large visual question answering dataset for cultural heritage. In: International Conference on Machine Learning, Optimization, and Data Science, pp. 193–197. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-95470-3\\_14](https://doi.org/10.1007/978-3-030-95470-3_14)
4. Barra, S., Bisogni, C., De Marsico, M., Ricciardi, S.: Visual question answering: which investigated applications? *Pattern Recogn. Lett.* **151**, 325–331 (2021)
5. Becattini, F., Ferracani, A., Landucci, L., Pezzatini, D., Uricchio, T., Del Bimbo, A.: Imaging Novecento. A mobile app for automatic recognition of artworks and transfer of artistic styles. In: Ioannides, M., et al. (eds.) EuroMed 2016. LNCS, vol. 10058, pp. 781–791. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-48496-9\\_62](https://doi.org/10.1007/978-3-319-48496-9_62)
6. Bongini, P., Becattini, F., Bagdanov, A.D., Del Bimbo, A.: Visual question answering for cultural heritage. In: IOP Conference Series: Materials Science and Engineering, vol. 949, p. 012074. IOP Publishing (2020)
7. Brown, T.B., et al.: Language models are few-shot learners (2020)
8. Cetinic, E., She, J.: Understanding and creating art with AI: review and outlook. *ACM Trans. Multimedia Comput. Commun. Appl. (TOMM)* **18**(2), 1–22 (2022)
9. Cornia, M., Stefanini, M., Baraldi, L., Cucchiara, R.: Meshed-memory transformer for image captioning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10578–10587 (2020)
10. Dale, R.: GPT-3: what’s it good for? *Nat. Lang. Eng.* **27**(1), 113–118 (2021)

11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
12. Elkins, K., Chun, J.: Can GPT-3 pass a writer's Turing test? *J. Cult. Analytics* **5**(2), 17212 (2020)
13. Fiorucci, M., Khoroshiltseva, M., Pontil, M., Traviglia, A., Del Bue, A., James, S.: Machine learning for cultural heritage: a survey. *Pattern Recogn. Lett.* **133**, 102–108 (2020)
14. Garcia, N., et al.: A dataset and baselines for visual question answering on art. In: Bartoli, A., Fusiello, A. (eds.) *ECCV 2020*. LNCS, vol. 12536, pp. 92–108. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-66096-3\\_8](https://doi.org/10.1007/978-3-030-66096-3_8)
15. Han, J., Zhang, D., Cheng, G., Liu, N., Xu, D.: Advanced deep-learning techniques for salient and category-specific object detection: a survey. *IEEE Signal Process. Mag.* **35**(1), 84–100 (2018)
16. Kheradpisheh, S.R., Ganjtabesh, M., Thorpe, S.J., Masquelier, T.: STDP-based spiking deep convolutional neural networks for object recognition. *Neural Netw.* **99**, 56–67 (2018)
17. Li, X., et al.: OSCAR: object-semantics aligned pre-training for vision-language tasks. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) *ECCV 2020*. LNCS, vol. 12375, pp. 121–137. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58577-8\\_8](https://doi.org/10.1007/978-3-030-58577-8_8)
18. Lin, C.Y.: ROUGE: a package for automatic evaluation of summaries. In: *Text Summarization Branches Out*, pp. 74–81 (2004)
19. Liu, S., Zhu, Z., Ye, N., Guadarrama, S., Murphy, K.: Improved image captioning via policy gradient optimization of spider. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 873–881 (2017)
20. Lu, J., Yang, J., Batra, D., Parikh, D.: Hierarchical question-image co-attention for visual question answering. In: *Advances in Neural Information Processing Systems*, pp. 289–297 (2016)
21. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318 (2002)
22. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training (2018)
23. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI Blog* **1**(8), 9 (2019)
24. Rajpurkar, P., Zhang, J., Lopyrev, K., Liang, P.: SQuAD: 100,000+ questions for machine comprehension of text. arXiv preprint [arXiv:1606.05250](https://arxiv.org/abs/1606.05250) (2016)
25. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems*, pp. 91–99 (2015)
26. Salton, G., Buckley, C.: Term-weighting approaches in automatic text retrieval. *Inf. Process. Manage.* **24**(5), 513–523 (1988)
27. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. arXiv preprint [arXiv:1910.01108](https://arxiv.org/abs/1910.01108) (2019)
28. Seidenari, L., Galteri, L., Bongini, P., Bertini, M., Del Bimbo, A.: Language based image quality assessment. In: *ACM Multimedia Asia*, pp. 1–7 (2021)
29. Sheng, S., Laenen, K., Moens, M.-F.: Can image captioning help passage retrieval in multimodal question answering? In: Azzopardi, L., Stein, B., Fuhr, N., Mayr, P., Hauff, C., Hiemstra, D. (eds.) *ECIR 2019*. LNCS, vol. 11438, pp. 94–101. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-15719-7\\_12](https://doi.org/10.1007/978-3-030-15719-7_12)

30. Shih, K.J., Singh, S., Hoiem, D.: Where to look: focus regions for visual question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4613–4621 (2016)
31. Stefanini, M., Cornia, M., Baraldi, L., Corsini, M., Cucchiara, R.: Artpedia: a new visual-semantic dataset with visual and contextual sentences in the artistic domain. In: Ricci, E., Rota Bulò, S., Snoek, C., Lanz, O., Messelodi, S., Sebe, N. (eds.) ICIAP 2019. LNCS, vol. 11752, pp. 729–740. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-30645-8\\_66](https://doi.org/10.1007/978-3-030-30645-8_66)
32. Su, W., et al.: VL-BERT: pre-training of generic visual-linguistic representations. arXiv preprint [arXiv:1908.08530](https://arxiv.org/abs/1908.08530) (2019)
33. Sun, T.X., Liu, X.Y., Qiu, X.P., Huang, X.J.: Paradigm shift in natural language processing. *Mach. Intell. Res.* **19**(3), 169–183 (2022)
34. Tan, H., Bansal, M.: LXMERT: learning cross-modality encoder representations from transformers. arXiv preprint [arXiv:1908.07490](https://arxiv.org/abs/1908.07490) (2019)
35. Vannoni, F., Bongini, P., Becattini, F., Bagdanov, A.D., Bimbo, A.: Data collection for contextual and visual question answering in the cultural heritage domain (2020)
36. Vaswani, A., et al.: Attention is all you need. In: Advances in Neural Information Processing Systems, pp. 5998–6008 (2017)
37. Vedantam, R., Lawrence Zitnick, C., Parikh, D.: CIDEr: consensus-based image description evaluation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4566–4575 (2015)
38. Wang, P., et al.: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. arXiv preprint [arXiv:2202.03052](https://arxiv.org/abs/2202.03052) (2022)
39. Yang, Z., et al.: An empirical study of GPT-3 for few-shot knowledge-based VQA. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, pp. 3081–3089 (2022)